International Institute for Applied Systems Analysis
I I A S A www.iiasa.ac.at

Address: IIASA, Schlossplatz 1, A-2361 Laxenburg, Austria

Email: repository@iiasa.ac.at **Department:** Advanced Systems Analysis | ASA

Working paper

Artificial Intelligence and Machine Learning for Systems Analysis of the 21st Century

Elena Rovenskaya (rovenska@iiasa.ac.at)
Keivan Aghababaei Samani (samani@cc.iut.ac.ir)
Artem Baklanov (baklanov@iiasa.ac.at)
Tatiana Ermolieva (ermol@iiasa.ac.at)
Christian Folberth (folberth@iiasa.ac.at)
Steffen Fritz (fritz@iiasa.ac.at)
Hadi Hadi (hadi@iiasa.ac.at)
Valeria Javalera-Rincon (javalera@iiasa.ac.at)
Andrey Krasovskii (krasov@iiasa.ac.at)
Finn Laurien (laurien@iiasa.ac.at)
Miguel Poblete Cazenave (poblete@iiasa.ac.at)
Thomas Schinko (schinko@iiasa.ac.at)
Mikhail Smilovic (smilovic@iiasa.ac.at)
Piotr Zebrowski (zebrowsk@iiasa.ac.at)

WP-19-010

Approved by:

Name Albert van Jaarsveld
Program: Director General and Chief Executive Officer
Date: 09 December 2019

Table of contents

ZVR 524808900

This research was funded by IIASA and its National Member Organizations in Africa, the Americas, Asia, and Europe.

Executive Summary

This paper overviews research being done at IIASA with use of machine learning (ML) methods based on the presentations and discussions at the internal workshop held at IIASA on June 24, 2019.

Machine learning (ML) has been developed in response to big data becoming available across various disciplinary domains. It treats an actual data-generating process as a black box, which takes values of predictor variables as an input and returns output in form of values of response variable(s) of interest. The field of ML is perceived as less formalized in comparison with classical statistics, but the theoretical foundations of ML are well developed and standard assumptions are less restrictive than those of classical statistics. ML aims at generalizing learning algorithms focusing on their performance in out-of-sample prediction.

There are several types of ML. The most prominent one is *supervised learning*, i.e., learning from samples containing both values of predictor variables (inputs) and values of response variables (outputs). Typical applications of supervised learning are classification, pattern recognition and regression-like problems. *Unsupervised learning* is a setting, in which available data contains output values without any information about corresponding inputs. The goal is to uncover interesting patterns in the data, such as clusters, latent factors, graph structures, correlations or directions of dependence. *Semi-supervised learning* is a mixture of the two above-mentioned learning settings, in which data contains information about inputs, but only some examples, on which an algorithm is trained, contain outputs of the black box. A notable example is image classification, where only a part of training images is labeled (i.e. interpreted by human). *Reinforcement learning* is another variation, in which inputs are known to a learning algorithm, but instead of full information about outputs, only some feedback on the performance of the algorithm is given. Most notable applications are optimization problems, in which the payoff landscape is not known to a learning agent a-priori – instead she explores it through trying different actions. Examples encompass optimization of performance of certain system operating in changing environment or playing a game, the rules of which are unknown.

Reinforcement learning and increasingly ML in general are often regarded as a part of Artificial Intelligence (AI), which studies intelligent agents, i.e. devices that are able to percept the surrounding environment and take rational actions in order to achieve predefined goals. Another rapidly expanding area in AI/ML is *deep learning*. In this paradigm, a learning agent is given a massive amount of highly dimensional training data (e.g. images) and autonomously discovers a multi-layer representation of data (from low-level to high-level features), which it then uses for prediction. This is in contrast to the traditional ML, in which features/predictors are given to the learning algorithm by the programmer. The workhorses of deep learning field are convolution and recurrent neural networks, which recently revolutionized image recognition and natural language processing.

It is important to realize what new opportunities AI/ML may offer to the applied systems analysis. In some cases, objectives of system analysis like prediction or optimization of system behavior can be achieved without having a process-based model of the system. For instance, reinforcement learning may generate robust strategies of managing a system which is not fully understood. Supervised learning can be used to emulate behavior of complex systems, making it easier to study their responses to forcing or external drivers. Data mining and unsupervised ML techniques may also help in making sense of massive data reflecting system's dynamic and its internal workings by uncovering patterns (e.g., clusters) and apparent relationships.

AI/ML methods should be used with understanding of their assumptions and limitations. Discovery of apparent patterns and relationships is only the first step. Next lays the challenge of interpreting outputs of learning algorithms and formulating theory explaining occurrence of discovered patterns – which is still beyond capabilities of existing AI systems.

Research carried out at IIASA using various methods from machine learning can be broadly categorized into three groups each comprising a few areas.

Data production and identification. IIASA runs crowdsourcing campaigns to map land cover based on data from satellite observations. Bad quality of images, as well as contradictory data from respondents hinder sufficient accuracy and robustness of identification. Image selection and vote aggregation procedures based on machine learning have been developed and tested at IIASA that enabled significant improvements of the identification quality. Text mining is used to understand people's sentiments. For example, it was used to assess the user perception and usability of a decision support tool on flood insurance developed at IIASA, which is very helpful to inform further development of this tool to make it more useful to users. Clusters (or communities using ecological language) in complex networked systems can be identified by means of machine learning based techniques.

Prediction of spatially detailed economic and environmental information. Random forest and other machine learning algorithms are used to build statistical models emulating data from observations or from process-based models in various areas of IIASA research. These include land cover maps for various countries and regions (e.g., Indonesia, Africa) at very fine resolutions; maps of net primary production of forests in Indonesia. Such statistical emulators allow for downscaling and producing data at fine resolutions using reasonable computational resources. Over time, the predictive power of such models steadily increases thanks to more data available and more advanced methods used. For example, a machine learning based emulator of IIASA's EPIC model enabled estimating crop yields in different policy scenarios in Mexico at fine resolutions. The large-scale Dynamic Global Vegetation Model LPJmL that simulates wheat yields was emulated using neural networks. The emulator was applied across northwestern India to evaluate different irrigation schedules. MODFLOW model was employed to simulate the amount of streamflow depletion as the result of groundwater pumping around Lake Michigan, USA. Its emulator constructed using ML was used to downscale model results to finer resolution levels. Using tree-based random forest algorithm to combine interview data and object-based data, IIASA researchers reconstructed spatially distributed loss distributions from floods in Germany. Reflecting on the power of machine learning in this case, first, we can state that a machine learning model can be a sufficiently good replacement of a process-based model; it can work faster and enable downscaling and transferability, for example, to data-challenging locations. Moreover, in some areas, for instance, in water, a machine learning model can be more universal (and, hence, transferrable) than a process-based model, which typically accounts for physical peculiarities of a region for which it has been developed. In a similar way, machine learning can help "simulate" data, i.e., extrapolate and interpolate as necessary. The gain in computing time comes at the cost of accuracy; yet, in many applications, for example, to explore the parameter or policy scenario space using machine learning based emulators is the only feasible way.

Policy optimization. To produce truly integrated solutions taking synergies and minimizing tradeoffs across sectors and regions, a reinforcement learning algorithm has been adapted by IIASA researchers to link individual and sectorial models into a single meta-model. This algorithm is currently being tested in a project, which aims to support international negotiations on SDGs. Importantly, the model linkage using reinforcement learning does not require sharing individual model codes and hence it can be organized flexibly in a decentralized manner; models can be added and removed, they can be of different nature – the only requirement is that their shared variables should be harmonized. Neural networks have been used to come up with a regional decision support tool to inform authorities on multi-dimensional health impacts of air quality policies. This decision support tool allows account for key nonlinearities in atmospheric phenomena, which are typically linearized in most state-of-the-art models, including IIASA's flagship GAINS model.

There are a few general challenges related to using ML-based models for understanding complex systems and for informing policy. As for any other model, one challenge is *model validation*. With more data available, out-of-sample prediction becomes a standard practice to validate statistical models, which also applies to ML-based models. But what is considered as a "good enough" prediction power varies from application to application. Another big conceptual challenge for using ML for systems analysis is that ML-based models are seen as "*black boxes*" by end-users. ML-based models are particularly difficult to communicate because they are often missing theoretical underpinnings and at the same time they are quite complex in terms of the choice of input variables. The concept of "explainable" AI is being developed now and IIASA might want to take advantage of it. Related to the black box issue is communication of uncertainty. ML-based models have the advantage of producing probabilistic estimates, yet it is known that communicating uncertainty ranges instead of a single number is a challenge. Theoretical limitations of ML lead to pitfalls like *false-positives*, *estimate biases*, and *limited power in dealing with structural changes*. *Computational resources* needed to run ML with big data is clearly another challenge. Related to this, ML-based models are *data-hungry*.

Systems analysis will surely take advantage of big data and ML techniques in the coming years. Among the promising new applications are comprehensive monitoring of SDGs, specifying uncertainties for robust decision support tools, merging qualitative and quantitative information – to name just a few. Newly available data such as commercial data (mobile phone records, financial data) or passive sources of data from citizens (social media, wearables) will open up new avenues.

Artificial Intelligence and Machine Learning for Systems Analysis of the 21st Century

Preamble

This paper overviews research being done at IIASA with use of machine learning (ML) methods. We elaborate on promising areas of application and advantages and challenges of using ML.

These reflections are done as a part of strategic planning process going on at IIASA at the moment, which aims to come up with a new research strategy for 2021-2030, as well as a supporting research plan. It has been recognized that while applications of ML in commercial sector are numerous and become more and more powerful day to day, it is not yet so common to use ML for creating societal impact.

To explore the opportunities in this context and to reflect on what IIASA's role might be, an internal working group was initiated. This paper emerged from the internal workshop held by the working group at IIASA on June 24, 2019; the workshop invited all IIASA scientists to contribute. The workshop program can be found in Appendix A to this paper.

1 What is machine learning (ML) and artificial intelligence (AI)?

Analysis of data carrying information about processes of interest is at the heart of many applied sciences, including systems analysis. In the 20th century, the field of data analysis took a great advantage of methods from classical statistics. In brief, the classical statistics' guiding paradigm can be described as a sequence of the following steps: (1) formulation of a hypothesis about the relation between variables describing the phenomenon to be analyzed; this hypothesis is then translated into a statistical model of the real process, through which nature generates the observed data; in this statistical model, the functional form of the relation between variables is specified; (2) obtaining of (preferably, unbiased) estimates of the statistical model's parameters; then testing whether the data indeed support the chosen hypothesis regarding the data generating process (by performing model's goodness-of-fit tests); (3) once confidence in the estimated statistical model has been established, it can be used for prediction. This process is also sometimes called a data modeling approach (see e.g., Breiman 2001a).

This approach has proven to be very useful in various applications from physics to economics to social sciences. However, with big data becoming more and more available in the 21st century across different disciplinary domains, its limitations related to high dimensionality of data are becoming increasingly ubiquitous. Statistical modeling of such data involves large numbers of predictor variables, which requires estimation of multiple parameters, reducing in turn the power of goodness-of-fit tests (i.e., increasing the probability of type-II errors). A standard response to this so-called "curse of dimensionality” is to reduce the set of predictor variables included in the model. This, however, may lead to a loss of important information; it also introduces a large degree of arbitrariness into the analysis, potentially undermining its results.

Another concern related to arbitrariness arises in situations when several statistical models (possibly of different types) happen to fit equally well to model data at hand. If a-priori knowledge about the nature of the phenomenon of interest is limited, it is difficult to justify an intelligent choice of a particular statistical model (for case examples, see Breiman 2001a). Any good statistical model provides information on prediction errors (e.g. in form of confidence intervals) – but under the assumption that this is a correct model to represent data. Yet, in situations where there are several competing models, an arbitrary choice of one of them introduces a new source of error: if the selected model is wrong, then also predictions made with use of this model may be wrong too. Classical statistics does not offer a way of accounting for and controlling this source of prediction errors.

The machine learning (ML) approach, known also as algorithmic modeling (Breiman 2001a), strives to control the prediction error resulting from use of inaccurate models. It treats an actual data-generating process as a black box, which takes values of predictor variables as an input and returns output in form of values of response variable(s) of interest. The ultimate resulting ML-model is rarely presented explicitly; contrary to classical statistics, modelers do not typically strive to give it a meaningful substantive interpretation.

In ML, it is assumed that internal mechanisms of this “nature’s black box” are too complex to be explicitly represented. Instead, ML aims at finding an algorithm mimicking the behavior of the black box that maximizes the predictive power (possibly at the cost of a bias). The field of ML is perceived as less formalized in comparison with classical statistics, but the theoretical foundations of ML are well developed and standard assumptions (that data points are independent and drawn from the same – possibly unknown – distribution) are less restrictive than those of classical statistics. ML aims at generalizing learning algorithms focusing on their performance in out-of-sample prediction. (This is in contrast to the classical statistics that is mainly concerned with testing how well statistical models fit the sample of data at hand.) The key element of ML is the relationship between the expected out-of-sample error, the complexity of algorithm, and the amount of available training data.

A standard procedure in ML is to train an algorithm on a subset of available data (training sample) and validate it by measuring its performance in out-of-sample prediction with use of the remaining data (testing sample). This is less formal than goodness-of-fit tests used in the classical statistics where every type of statistical models requires customized and often technically advanced goodness-of-fit tests. Another difference between the ML approach and the classical statistics is that ML methods suffer from the curse of dimensionality to a lesser degree. Indeed, ML algorithms are computationally more efficient and typically have better generalization properties, allowing to extract more information from high dimensional data than methods of classical statistics.

There are several types of ML (see e.g. Murphy 2012). The most prominent one is *supervised learning*, i.e., learning from samples containing both values of predictor variables (called features) and values of response variables, that is both the input and the output of the “nature’s black box”. Typical applications of supervised learning are, for instance, classification, pattern recognition and regression-like problems.

Unsupervised learning is a setting, in which available data contains values of outputs of the black box (representing data generating process) without any information about corresponding inputs. The goal of unsupervised learning is to uncover “interesting patterns” in the data, such as clusters, latent factors, graph structures, correlations or directions of dependence. Unsupervised ML overlaps with fields of “knowledge discovery” and data mining (there is no clear-cut delineation between them).

Semi-supervised learning is a mixture of the two above-mentioned learning settings, in which data contains information about inputs, but only some examples, on which an algorithm is trained, contain outputs of the black box. A notable example of the semi-supervised learning application is image classification where only a (usually small) part of training images is labeled (i.e. interpreted by human).

Reinforcement learning is another variation, in which inputs are known to a learning algorithm, but instead of full information about outputs, only some feedback on the performance of the algorithm is given. Most notable applications of reinforcement learning are optimization problems, in which payoff landscape is not known to a learning agent a-priori; instead she explores it through trying different actions. Examples encompass optimization of performance of certain system operating in changing environment or playing a game, the rules of which are unknown.

Reinforcement learning (and increasingly ML in general) is often regarded as a part of a vibrant discipline of Artificial Intelligence (AI), which studies intelligent agents, i.e. devices that are able to percept the surrounding environment and take rational actions (optimal in some sense, given the available knowledge about environment) in order to achieve predefined goals (see e.g. Russel and Norvig 2010).

Another rapidly expanding subfield of AI/ML is the so-called *deep learning* (see e.g. LeCun et al. 2015). In this paradigm a learning agent is given a massive amount of highly dimensional training data (e.g. images) and autonomously discovers a multi-layer representation of data (from low-level to high-level features), which it then uses for prediction. This is in contrast to the traditional ML, in which features/predictors are given to the learning algorithm by the programmer. The workhorses of deep learning field are convolution and recurrent neural networks, which recently revolutionized image recognition and natural language processing.

It is important to realize what new opportunities AI/ML may offer to the applied systems analysis. The bulk of applied systems modeling has much in common with the approach used in classical statistics, with much effort being put into detailed representation of systems’ elements and their interactions with the hope that this will lead to accurate predictions and/or identification of optimal management policies. This is similar to the approach of classical statistics focused on getting a “right” model for data generating process (which in turn enables prediction). As a result, systems modeling and statistical data modeling suffer from similar problems, e.g., over-parametrization and challenges of data integration, or large degree of arbitrariness in assumptions in face of knowledge gaps – to name just a few. Yet, the objectives of system analysis like prediction/optimization of systems’ behavior may be achieved without an exact and detailed model representation of the system of interest – similarly to AI/ML approach to prediction without explicit representation of data generating process. Here AI/ML appear to present alternative/complementary ways forward for systems analysis.

For instance, reinforcement learning may shed light on robust strategies of managing a system which is not fully understood (as is the case with, for example, coupled environment–economy systems). Supervised learning can be used to emulate behavior of complex systems, making it easier to study their responses to forcing or external drivers. Data mining and unsupervised ML techniques may also help in making sense of massive data reflecting system’s dynamic and its internal workings by uncovering patterns (e.g., clusters) and apparent relationships. In the same vein, deep learning may suggest most informative features, at least from the perspective of achieving certain objectives. Thus, AI/ML methods have potential to support exploratory analysis and formulation of research hypotheses.

AI/ML methods should be used with caution, however, and with understanding of their assumptions and limitations. Discovery of apparent patterns and relationships is only the first step. It is then up to researchers to verify whether they indeed hold true in reality or are they only artifacts of methods employed in the analysis. Next lays the challenge of interpreting outputs of learning algorithms and formulating theory explaining occurrence of discovered patterns – which is still beyond capabilities of existing AI systems.

In the next section, we overview examples of IIASA research using ML methods in a range of applications.

2 Machine learning in systems analysis at IIASA up to now

2.1 Inference of information from crowdsourcing

It is a new global trend to improve existing land cover monitoring tools to involve general public in scientific research (Schepaschenko et al. 2019). More specifically, crowdsourcing campaigns are conducted to identify different land cover types in locations based on satellite images of these locations. The challenge here is to aggregate multiple, possibly contradictory identification results from different human subjects on the same image and come up with an accurate final identifier. To address this challenge, a vote-aggregation procedure was developed by combining state-of-the-art ML algorithms (RandomForest, Xgboost) and some additional heuristics. In the case study of the Cropland Capture Game (Salk et al. 2016), the application of this developed procedure enabled an increase of the estimated consistency with expert opinions from 77 to 96% (Baklanov et al. 2016; Baklanov et al. 2017).

Low quality of images is a serious obstacle for identification by human subjects. One way to tackle this problem is to remove blurry images that can be confusing from the training set and not send them to human subjects for identification. Baklanov et al. applied threshold-based blur detection algorithms for such a removal, which led to increase of the estimated accuracy of simple majority voting from 77% to 90% (Baklanov et al. 2016).

2.2 Construction of maps based on satellite observations

Creating maps depicting various geographically explicit data obtained based on satellite images is an area of systems analysis, where supervised machine learning has been the major workhorse since a few decades. Typically, data to be depicted on maps are represented as a discrete variable, for example, it can be land cover types such as bare soil, grass cover, shrub cover, tree cover, water and artificial surfaces. In some cases, it can also be a continuous variable, for example, net primary production (NPP) of forests or economic parameters, such as the GDP or population income per capita.

Google made available a map production system called Google Earth Engine with some classification algorithms, which can now be used to generate new maps for next years to come. The European Copernicus program with its Sentinel 1 and 2 satellites is providing useful consistent high-resolution (10 m) data for generation of new maps of time series. In particular in the future it is anticipated that much better information on land use and land cover change will become available.

Mathematically, it is a pattern recognition or a classification problem, where input data include sets of multiple satellite images of the same locations. Despite having several images of the same place available, the difficulty here is in the quality of images, for example, while we are interested in what is on the ground, images are often dominated by clouds.

Inputs to an ML algorithm sometimes include also other auxiliary data, which can act as additional predictors to improve the quality of prediction. To inform the training stage of a supervised machine learning algorithm, a campaign, in which images from the training set are identified by humans, needs to be run. ML is known to be "hungry for data", hence these campaigns ideally involves thousand people. In 2000s, IIASA pioneered an effective and efficient tool for collecting such data called GeoWiki (Fritz et al. 2009). The current number of actively registered users is 1500.

As for machine learning algorithms used in this area, historically, they ranged from supervised classification to decision trees, to initial ML approaches, which started with the random forest and now moved to convolutional neural networks. Each newly used family of methods was able to ensure better predictive power as tested via out-of-sample predictions (Fritz 2019).

Thanks to the EU investments in its satellite infrastructure and to other initiatives, data currently publicly available for training and testing of ML algorithms is really massive, which makes the prediction quality generally satisfactory.

In case of land use applications in particular, for these maps to be maximally useful to inform policy, they should have a sufficiently fine resolution, as decisions are often taken at the level of municipalities. The community is working on moving to finer-resolution maps. One challenge here is that auxiliary information used to ensure satisfactory levels of predictive power is sometimes available at a coarser resolution. For example, this is the case for climate and precipitation data. Resolution (in)compatibility is also a challenge in the temporal dimension. Finer resolutions are particularly important for mountainous areas where vegetation types change quickly as climate conditions vary with changing altitude.

IIASA is actively working on making land use and vegetation type maps available, also as open access. For example, IIASA scientists in collaboration with local experts are constructing annual land cover maps for Indonesia for 2015-2018 at the resolution of 30 m using a random forest algorithm (Hadi et al. 2019). Google earth engine is used as a platform here to undertake the massive processing of large amounts of satellite data. With the currently existing preliminary results, the overall accuracy in this case is 45% (this validation was carried out with reference data which was collected via visual interpretation of very high resolution google earth and ground data). This relatively low accuracy is due to large errors in certain parts of Indonesia, which requires additional attention and data preparation. This will be tackled as a next step in order to reach a significantly higher accuracy. A land cover map for the entire African continent was constructed for the year 2015 at 20 m resolution by the ESA CCI land cover consortium (Fritz 2019). The overall accuracy here reached 65%. Another example of such work is the work of Lesiv et al., who contributed to the Copernicus Global Land Service[1] by making available global land cover maps at 100 m resolution.

Machine learning was used to estimate net primary productivity (NPP) of planted forests in Indonesia (Krasovskii et al. 2019). Remote sensing data were combined with soil information, a range of climate data including temperatures, precipitation, relative humidity etc., as well as land cover map. In this study, scientists explicitly tested several ML methods including KNeighbors regression, support vector regression, Gaussian process regression, and random forest trees, and found out the latter delivering the highest prediction quality in terms of the mean absolute percentage error (96%). The random forest tree algorithm allowed ranging the explanatory biophysical parameters driving the forest dynamics according to their importance, in particular, soil water turned out to be most important predictor of the NPP in this case study. Soil water comes from the water balance routine taking into account soil texture and potential evapotranspiration calculated by the Hargreaves method, which takes into account mean monthly radiation, as well as monthly minimum and maximum temperature. Coupled with the statistical analysis of the growth curves, ML helped identifying spatially explicit site index for several

[1] https://land.copernicus.eu/global/products/lc

plantation species in Indonesia. The method was used to find optimal rotation times in IIASA's global spatially explicit forestry model G4M. Practical examples for plantation forestry in Indonesia (including acacia and teak species) were demonstrated (Krasovskii et al. 2019).

With more and more satellite data becoming available, the challenge of data storage becomes more and more severe despite progress in these technologies. Data needs challenged the capacity limits, for example, in the project on mapping palm oil plantations (Baklanov et al. 2018; Baklanov et al. 2019b). In this project on monitoring industrial oil palm plantations (OPPs), a novel approach employing a Fully Convolutional Neural Networks (FCNs) to solve Semantic Segmentation Problem for Landsat imagery was used. The proposed approach significantly outperforms state-of-the-art per-pixel classification methods (Baklanov et al. 2019a). Moreover, the trained FCN is robust to spatial and temporal shifts of input data. This shows that FCNs enable OPPs mapping of entire countries and may serve for yearly detection of oil palm expansion. Data requirements were also a limiting factor in another project, in which an African land cover map was constructed: here the data used was as much as 465 TB (Fritz 2019).

2.3 Model emulation and downscaling

The EPIC (Environmental Policy Integrated Climate; Williams 1990) field-scale model is a complex agronomic model for assessments of agricultural management and climate impacts on crop yields and the environment. The spatially explicit, global EPIC-IIASA model is an example of a global policy model developed and used by IIASA to provide advice to policy makers. EPIC-IIASA enables assessing how land management affects the environment. It can be used to compare land and forest management systems and their effects on environmental indicators like water availability, nitrogen and phosphorus levels in soil, and greenhouse gas emissions. Large-scale crop models like EPIC-IIASA are also important tools for agricultural climate change impact assessments, as input data provider for land use modelling, and most recently for farmer information.

EPIC-IIASA is a global model operating at 50 km resolution (appr. 35,000 grid cells globally). This resolution is however not enough to inform policy at the local level, hence the modelers introduced an "enhanced" grid of 1 km resolution. There is therefore a need to downscale the model results to this resolution level and Folberth et al. (2019b) developed a ML-based approach to do so. Using the original crop model itself, such a downscaling would require a wide range of soil, climate, and management variables in specific formats for each target region and management system. This results in long lead times to obtain regionally or location specific information.

Hence, while the low data availability and high demand for computational resources are two obstacles for straightforward downscaling (Baklanov et al. 2019a), ML was used to emulate the EPIC model; namely, a ML algorithm was trained on the coarser resolution and then it was applied to derive predictions on the finer resolution level. To inform this downscaling, auxiliary data including soil, climate, and site data were used. As the data on these auxiliary variables are available at the finer resolution, the emulator can effectively predict the environmental impacts of land management policies at the desired resolution levels. So far, this model emulation was done as a proof-of-concept, which is why 80% of the available data were used for training and the remaining 20% were used for testing for out-of-sample prediction power. In this exercise, in case of 50 km resolution, R2 reached as high as 99% at the global level. Essential covariates are sparser than those required by the crop model and can be provided in generic format, rendering the retrieval of crop yield predictions highly efficient. In the case study of Mexico, whose data was not a part of the training dataset, R2 reached as high as 97%. In this study, two ML methods were compared: random forests and extreme gradient boosting. While both delivered similar predictive powers in terms of correlation with simulated model outputs, extreme gradient boosting reproduced more accurately also local inter-annual yield variability and required much less (~8 times less) computational time on the same computational infrastructure (Folberth et al. 2019a).

Model emulation of EPIC is, therefore, useful for swift estimation of crop yields in areas, where observations are insufficient. Crop yield predictions obtained via our machine learning approach can be operationalized for farmer information, e.g. in an app recently developed to provide location-specific yield potentials. Presently based on static simulation outputs, these can be updated on-the-fly based on farmers' input data on crop management and soil characteristics among others. Further potential applications include rapid screening of management alternatives and identification of regional hotspots for interventions in policy-making and project planning (Folberth et al. 2019a). After this proof of concept, additional dimensions on crop management and an extended range of climate conditions will be included. Further exploration will cover data assimilation of crop model outputs and field observations.

However, downscaling results suffer from the "extrapolation curse" of ML: in case we would like to predict effects under policy scenarios not included in the training dataset (or are a small portion of the entire dataset), the prediction outcomes will likely to be much less accurate (Folberth et al. 2019a). To avoid this, the training data will be expanded to cover a wider range of climates and a testing routine will be implemented to evaluate *a priori* whether predictors are within the distributions of training data and the application of the ML algorithm can be considered safe.

In another study, IIASA researchers used an emulator of a crop yield model to estimate effects of different irrigation schedules. The large-scale Dynamic Global Vegetation Model LPJmL that simulates wheat yields was emulated using neural networks. The emulator was applied across northwestern India under ~20,000 different irrigation schedules (Smilovic et al. 2019).

In the area of water, MODFLOW model was employed to simulate the amount of streamflow depletion as the result of groundwater pumping around Lake Michigan, USA. Its emulator constructed using ML was used to downscale model results to finer resolution levels. In-sample prediction quality of this emulator was 85%.

Model emulation is a promising direction for policy models to become easier usable at their original, as well as finer resolutions. Model emulation allows to approximate the results of the process-based models while decreasing the computational intensity and potentially allowing for predictions in areas with similar characteristics.

2.4 Estimation of loss distributions from natural disasters

Loss distributions from natural disasters are known to be non-Gaussian, particularly at lower geographical scales or at the object level. Therefore, point estimates (e.g., means) are not very meaningful to inform decisions on insurance and ex-ante mitigation measures. It is a big methodological challenge in the area of flood risk assessment to properly account for major uncertainties spreading across the three components of hydrometeorological risk: hazards, exposures, and vulnerabilities, and to do that consistently across spatial scales. The state-of-the-art approaches have limited capabilities in this regard, which often leads to under- or overestimation of risks.

In one IIASA study based on a YSSP project (Sieg et al. 2019), scientists used a tree-based random forest algorithm (Breiman 2001b; Hothorn et al. 2006) to combine interview data and object-based exposure data from OpenStreetmap.org in order to estimate direct riverine flood damages in the case study of the federal state of Saxony, Germany. The associated uncertainties of input variables describing hazard and exposure, and the uncertainty within the damage estimation are modeled by probability distributions. Interviews, which were held with 557 companies after the 2002 and 2013 flood events, resulted in a dataset of materialized direct damages at the company level that was used to train the random forest model. Data from OpenStreetmap.org included specific information such as the occupancy or height of buildings at the object level. The random forest model enabled to reconstruct the entire direct loss distribution for the flood event in 2013 in Germany. The latter was done by sampling 1,000 possible damage realizations for each affected company, grouped into 19 economic sectors. The obtained post-event estimates were shown to be more accurate than those derived from land use – based top-down models, which are the common workhorse in this area. The developed approach is therefore suitable for the analysis of future risks with the possibility of including any kind of uncertainty (Schinko et al. 2019).

2.5 Identification of communities in networks

Communities are sub-graphs of a network whose nodes are densely connected to each other, but there are much fewer connections between different sub-graphs. Information on communities may be useful for simplification of network models; it can also be a basis for tailoring policy interventions. Machine learning offers tools to identify communities called matrix factorization methods.

Non-negative Matrix Factorization (NMF) method was used to detect a given number of community structures in complex networks. The method requires an iterative procedure, which typically converges very fast. The method is successfully applied to a number of real-world and computer-generated networks, including protein-protein interaction networks, sport club networks dolphin social networks (Zarei et al. 2009). The approach can be used for binary, as well as for weighted networks (Aghababaei Samani 2019).

2.6 Decision support systems based on neural networks

IIASA is using a number of decision support models, which generate normative scenarios by evaluating alternative solutions and choosing the most cost-effective one. Mathematically, these models are formulated as linear programming (LP) problems.

For example, IIASA's GAINS model is a decision support model that was used for decades to inform Environmental Authorities to plan air quality policies that fulfill e.g., EU Directive requirements in a cost-efficient way. In this model, the air quality indices are computed using linear models, derived through model reduction techniques starting from deterministic Chemical Transport Model simulations. This linear approach limits the applicability of these surrogate models (Wagner, abstract). Wagner et al. used a multi-objective nonlinear approach to control air quality at a regional scale. Both economic and air quality sides of the problem were modeled through artificial neural network models. Simulating the complex nonlinear atmospheric phenomena, they can be used in an optimization routine to identify efficient solutions of a decision problem for air quality planning. The methodology was applied over Northern Italy, an area in Europe known for its high concentrations of particulate matter (Wagner, abstract).

2.7 Model linkage using reinforcement learning

Producing integrated solutions delivering on multiple dimensions of sustainability requires integrated modeling, that is linking sectorial and regional models together respecting joint resource constraints. While in theory "hard" model linkage is the best way to have a truly coherent and self-consistent multi-sector multi-region model, in reality reprogramming sectorial and regional models into a single code is tedious task, which is also often not feasible because models are owned by separate teams and are implemented in different environments without necessarily full documentation.

Instead of hard model linkage, Javalera- Rincón proposed using a reinforcement learning algorithm to organize a distributed model linkage via a "central hub" called "linker agent". In the training phase, the linker agent, in a trial-end-error fashion, learns from individual models by exploring the action (solution) space and collecting rewards. The proposed algorithm is based on the standard Q-learning approach and has two different versions called planningByInstruction and planningByExploration (Javalera-Rincón 2019b). In the exploitation phase, accumulated rewards are used to first prioritize feasible actions, and then among feasible actions prioritize those, which deliver a socially optimal solution (Javalera-Rincón 2019a).

2.8 Text mining

In the current realities of the digital age, social scientists are confronted with increasing challenges of monitoring and assessing larger and more complex sets of data collected by qualitative or semi-quantitative methods. Semantic text analysis methods based on machine learning algorithms can guide social scientists to better understand, structure, cluster and assess the large amount of data.

Laurien et al. applied a semantic text analysis method to more than 30,000 user feedbacks and comments to assess the user perception and the usability of the Flood resilience measurement tool. The approach is able to assign grades to user statements to describe statement sentiments (positive, neutral, negative). The approach is able to highlight strengths and weaknesses for specific measures collecting complex system elements of community flood resilience (Laurien et al. 2019).

Rekabsaz et al. used Information Retrieval (IR) term weighting models extended by related terms using word embeddings to predict financial market volatility. They relied on a fusion of sentiment analysis of annual disclosures of companies in stock markets and some selected market data (Rekabsaz et al. 2017).

2.9 Alternative approaches to error modeling

The entire classical statistics and subsequently ML is based on the assumption that the mean squared error is the primary metric to evaluate the goodness-of-fit of a prediction and, therefore, algorithms are arranged in a way to minimize this metric. The mean squared error (MSE) has convenient properties of being a smooth function, which reasonably disfavors outliers. Under the assumption that data errors are subject to a Gaussian distribution, various useful statistics (such as P-value etc.) are computed to describe the power of a relationship estimated.

However, MSE-based estimates may be sensitive to outliers, especially when the errors are non-Gaussian. MSE-based models are capable of predicting an average response, but not necessarily potential extreme outcomes. Therefore, in some applications, it may make sense to use other metrics for the goodness-of-fit, which take into consideration the spread of the response variable, such as, for example, a quantile loss (QL) function (Ermolieva et al. 2019). Due to their nature, quantile-based estimates are often more robust with respect to outliers than MSE-based predictions (see Ermoliev and Hordijk 2006 for different facets of robustness of the QL function).

Quantile-based functions are non-smooth, which poses a challenge for using gradient decent methods. Instead, it is proposed to apply iterative quasi-gradient procedures to derive robust parameter estimates (Ermoliev 2009). Such estimates were used to estimate crop yields distributions, which are often negatively skewed because of biological constraints that limit plants growth in response to various, often cumulative, combinations of weather parameters (precipitation, temperature, pressure, etc.).

3 Some conclusions and potential outlook

This document summarizes research carried out at IIASA using various methods from machine learning. This overview is certainly incomplete, but we feel that it covers major directions and case studies, which can be broadly categorized into three groups each comprising a few areas.

Data production and identification. IIASA runs crowdsourcing campaigns to map land cover based on data from satellite observations. Bad quality of images, as well as contradictory data from respondents hinder sufficient accuracy and robustness of identification. Image selection and vote aggregation procedures based on machine learning have been developed and tested at IIASA that enabled significant improvements of the identification quality (see section 2.1). Text mining is used to understand people's sentiments. For example, it was used to assess the user perception and usability of a decision support tool on flood insurance developed at IIASA, which is very helpful to inform further development of this tool to make it more useful to users (see section 2.7). Clusters (or communities using ecological language) in complex networked systems can be identified by means of machine learning based techniques (see section 2.4).

Prediction of spatially detailed economic and environmental information. Random forest and other machine learning algorithms are used to build statistical models emulating data from observations or from process-based models in various areas of IIASA research. These include land cover maps for various countries and regions (e.g., Indonesia, Africa) at very fine resolutions; maps of net primary production of forests in Indonesia (see section 2.2). Such statistical emulators allow for downscaling and producing data at fine resolutions using reasonable computational resources. Over time, the predictive power of such models steadily increases thanks to more data available and more advanced methods used. For example, a machine learning based emulator of IIASA's EPIC model enabled estimating crop yields in different policy scenarios in Mexico at fine resolutions. The large-scale Dynamic Global Vegetation Model LPJmL that simulates wheat yields was emulated using neural networks. The emulator was applied across northwestern India to evaluate different irrigation schedules. MODFLOW model was employed to simulate the amount of streamflow depletion as the result of groundwater pumping around Lake Michigan, USA. Its emulator constructed using ML was used to downscale model results to finer resolution levels (see section 2.2). Using tree-based random forest algorithm to combine interview data and object-based data, IIASA researchers reconstructed spatially distributed loss distributions from floods in Germany (see section 2.3). Reflecting on the power of machine learning in this case, first, we can state that a machine learning model can be a sufficiently good replacement of a process-based model; it can work faster and enable downscaling and transferability, for example, to data-challenging locations. Moreover, in some areas, for instance, in water, a machine learning model can be more universal (and, hence, transferrable) than a process-based model, which typically accounts for physical peculiarities of a region for which it has been developed. In a similar way, machine learning can help "simulate" data, i.e., extrapolate and interpolate as necessary. The gain in computing time comes at the cost of accuracy; yet, in many applications, for example, to explore the parameter or policy scenario space using machine learning based emulators is the only feasible way.

Policy optimization. To produce truly integrated solutions taking synergies and minimizing tradeoffs across sectors and regions, a reinforcement learning algorithm has been adapted by IIASA researchers to link individual and sectorial models into a single meta-model. This algorithm is currently being tested in a project, which aims to support international negotiations on SDGs (see section 2.6). Importantly, the model linkage using reinforcement learning does not require sharing individual model codes and hence it can be organized flexibly in a decentralized manner; models can be added and removed, they can be of different nature – the only requirement is that their shared variables should be harmonized. Neural networks have been used to come up with a regional decision support tool to inform authorities on multi-dimensional health impacts of air quality policies. This decision support tool allows account for key nonlinearities in atmospheric phenomena, which are typically linearized in most state-of-the-art models, including IIASA's flagship GAINS model (see section 2.5).

There are a few general challenges related to using ML-based models for understanding complex systems and for informing policy. As for any other model, one challenge here is *model validation*. With more data available, out-of-sample prediction becomes a standard practice to validate statistical models, which also applies to ML-based models. But what is considered as a "good enough" prediction power varies from application to application. For example, Hadi et al. (2019) consider 45% accuracy in out-sample-prediction of the annual land cover maps for Indonesia a good achievement. On the other hand, 97% accuracy was achieved in the case study predicting crop yields in Mexico based on EPIC model (Folberth et al. 2019a).

Another big conceptual challenge for using ML for systems analysis is that ML-based models are seen as "*black boxes*" by end-users. It is again not a unique challenge for ML-based models only. Developers of any more or less complex model have a challenge to communicate the model, its foundations, major assumptions and limitations. ML-based models are particularly difficult to communicate because they are often missing theoretical underpinnings and at the same time they are quite complex in terms of the choice of input variables. The concept of "explainable" AI (Goebel R et al. 2018) is being developed now and IIASA might want to take advantage of it. Related to the black box issue is communication of uncertainty. ML-based models have the advantage of producing probabilistic estimates, yet it is known that communicating uncertainty ranges instead of a single number is a challenge.

Theoretical limitations of ML lead to pitfalls like *false-positives* (Riley 2019), *estimate biases* (Poblete-Cazenave 2019), and *limited power in dealing with structural changes* (Poblete-Cazenave 2019). Prediction power outside of the range of the available data for policy parameters is often problematic, so the analysts run into the extrapolation problem.

Computational resources needed to run ML with big data is clearly another challenge. Related to this, ML-based models are *data-hungry*. In the area of land cover and vegetation mapping, crowdsourcing and satellite data is big and allows running advanced ML algorithms with satisfactory prediction power. Similarly, big data describing other human activities, also coming from satellites, can be used to produce maps detailing other indicators useful to policy makers. For example, in the energy area, in one YSSP project a map describing multi-dimensional energy poverty in sub-Saharan Africa will be constructed based on remote sensing data on nightlights (Poblete Cazenave 2019). Object-level data available enable estimating expose and vulnerability of buildings and infrastructure to natural disasters, which allows for accurate estimating of loss distributions. This information can then be used in risk analysis and to support decisions on appropriate ex-ante risk mitigation measures. Even if available data is big, there is a challenge of its accuracy. Detecting outliers in terms of inaccuracy and dealing with them in the best possible way requires ML-based approaches in itself. The minimal requirements to how big the data should be for a ML-based model to be reliable are not clearly defined. Some participants of the workshop were optimistic that at least with time, ML will not necessarily require big data, while some were skeptical about that.

Systems analysis will surely take advantage of big data and ML techniques in the coming years. Among the promising new applications are comprehensive monitoring of SDGs (Fritz, presentation), specifying uncertainties for robust decision support tools, merging qualitative and quantitative information – to name just a few. Newly available data such as commercial data (mobile phone records, financial data) or passive sources of data from citizens (social media, wearables) will open up new avenues.

References

Aghababaei Samani K (2019). Detecting community structure of complex networks by Non-negative Matrix Factorization (NMF) method. In: Artificial Intelligence and Machine Learning for Systems Analysis of the 21st Century, 24 June 2019, Laxenburg, Austria.

Baklanov A et al. (2019a). Applications of machine learning to inform environmental decision making. In: Artificial Intelligence and Machine Learning for Systems Analysis of the 21st Century, 24 June 2019, Laxenburg, Austria.

Baklanov A, Fritz S, Khachay M, Nurmukhametov O, Salk C, See L & Shchepashchenko D (2016). Improved vote aggregation techniques for the Geo-Wiki cropland capture crowdsourcing game. In: European Geosciences Union (EGU) General Assembly 2016, 17–22 April 2016, Vienna, Austria.

Baklanov A, Fritz S, Khachay M, Nurmukhametov O, Salk C, See L & Shchepashchenko D (2017). Vote aggregation techniques in the Geo-Wiki crowdsourcing game: A case study. In: Ignatov D et al. (eds) Analysis of Images, Social Networks and Texts. AIST 2016. Communications in Computer and Information Science 661. Springer, Cham, Switzerland.

Baklanov A, Khachay M & Pasynkov M (2018). Application of fully convolutional neural networks to mapping industrial oil palm plantations. In: van der Aalst W et al. (eds) Analysis of Images, Social Networks and Texts. AIST 2018. Lecture Notes in Computer Science 11179. Springer, Cham, Switzerland.

Baklanov A, Khachay M & Pasynkov M (2019b). Fully convolutional neural networks for mapping oil palm plantations in Kalimantan. In: Battiti R, Brunato M, Kotsireas I & Pardalos P (eds) Learning and Intelligent Optimization. LION 12 2018. Lecture Notes in Computer Science 11353. Springer, Cham, Switzerland.

Breiman L (2001a). Statistical modeling: The two cultures. Statistical Science 16: 199–231.

Breiman L (2001b). Random forests. Machine Learning 45: 5–32.

Ermoliev Y (2009). Stochastic quasigradient methods. In: Floudas CA & Pardalos PM (eds) Encyclopedia of Optimization, Springer, Boston, USA.

Ermoliev Y & Hordijk L (2006). Global changes: Facets of robust decisions. IIASA Interim Report. IIASA, Laxenburg, Austria: IR-06-001.

Ermolieva T et al. (2019). Artificial Intelligence: Non-smooth stochastic optimization and iterative quasigradient procedures for machine learning. In: Artificial Intelligence and Machine Learning for Systems Analysis of the 21st Century, 24 June 2019, Laxenburg, Austria.

Folberth C et al. (2019a). Combining crop modelling and machine learning for rapid provision of high-resolution crop yield predictions. In: Artificial Intelligence and Machine Learning for Systems Analysis of the 21st Century, 24 June 2019, Laxenburg, Austria.

Folberth C, Baklanov A, Balkovic J, Skalsky R, Khabarov N & Obersteiner M (2019b). Spatio-temporal downscaling of gridded crop model yield estimates based on machine learning. Agricultural and Forest Meteorology 264: 1–15.

Fritz S (2019). AI activities in EOCS. In: Artificial Intelligence and Machine Learning for Systems Analysis of the 21st Century, 24 June 2019, Laxenburg, Austria.

Fritz S, McCallum I, Schill C et al. (2009). Geo-Wiki.Org: The use of crowdsourcing to improve global land cover. Remote Sensing 1: 345–354.

Goebel R et al. (2018). Explainable AI: The new 42? In: Holzinger A, Kieseberg P, Tjoa A & Weippl E (eds) Machine Learning and Knowledge Extraction. CD-MAKE 2018. Lecture Notes in Computer Science, 11015. Springer, Cham, Switzerland.

Hadi et al. (2019). Mapping land cover in Indonesia at national scale using supervised machine learning algorithms and cloud computing. In: Artificial Intelligence and Machine Learning for Systems Analysis of the 21st Century, 24 June 2019, Laxenburg, Austria.

Hothorn T, Hornik K & Zeileis A (2006). Unbiased recursive partitioning: A conditional inference framework. Journal of Computational and Graphical Statistics 15: 651–674.

Javalera-Rincón V et al. (2019a). Reinforcement Learning Approach for Cooperative Control of Multi-Agent Systems. In: Artificial Intelligence and Machine Learning for Systems Analysis of the 21st Century, 24 June 2019, Laxenburg, Austria.

Javalera-Rincón V, Cayuela VP, Seix BM & Orduña-Cabrera F (2019b). Reinforcement learning approach for cooperative control of multi-agent systems. In: Proceedings of the 11th International Conference on Agents and Artificial Intelligence (ICAART 2019) pp 80–91, Porto, Portugal.

Krasovskii A et al. (2019). Application of machine learning to modeling planted forest productivity. In: Artificial Intelligence and Machine Learning for Systems Analysis of the 21st Century, 24 June 2019, Laxenburg, Austria.

Laurien A et al. (2019). Making smarter and faster decisions based on machine learning. In: Artificial Intelligence and Machine Learning for Systems Analysis of the 21st Century, 24 June 2019, Laxenburg, Austria.

LeCun Y, Bengio Y & Hinton G (2015). Deep learning. Nature 521: 436–444.

Murphy KP (2012). Machine Learning: A Probabilistic Perspective. MIT Press, Cambridge, USA.

Poblete Cazenave M (2019). The limits of inference without theory. In: Artificial Intelligence and Machine Learning for Systems Analysis of the 21st Century, 24 June 2019, Laxenburg, Austria.

Rekabsaz N, Lupu M, Baklanov A, Hanbury A, Duer A & Anderson L (2017). Volatility prediction using financial disclosures sentiments with word embedding-based IR models. In: Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL 2017) pp 1712–1721, Vancouver, Canada.

Riley P (2019). Three pitfalls to avoid in machine learning. Nature 572: 27–29.

Russell S & Norvig P (2010). Artificial Intelligence: A Modern Approach. 3rd Edition, Prentice-Hall, Upper Saddle River, USA.

Salk CF, Sturn T, See L, Fritz S & Perger C (2016). Assessing quality of volunteer crowdsourcing contributions: Lessons from the Cropland Capture game. International Journal of Digital Earth 9: 410–426.

Schepaschenko D, See L, Lesiv M, Bastin JF, Mollicone D, Tsendbazar NE, Bastin L, McCallum I, Bayas JCL, Baklanov A, Perger C, Dürauer M & Fritz S (2019). Recent advances in forest observation with visual interpretation of very high-resolution imagery. Surveys in Geophysics 40: 839–862

Schinko T et al. (2019). Integrated assessment of short-term direct and indirect economic flood impacts including uncertainty quantification. In: Artificial Intelligence and Machine Learning for Systems Analysis of the 21st Century, 24 June 2019, Laxenburg, Austria.

Sieg T, Schinko T, Vogel K, Mechler R, Merz B & Kreibich H (2019). Integrated assessment of short-term direct and indirect economic flood impacts including uncertainty quantification. PLoS ONE 14: e0212932.

Smilovic M et al. (2019). A tale of two studies: Neural networks for irrigation scheduling and streamflow depletion. In: Artificial Intelligence and Machine Learning for Systems Analysis of the 21st Century, 24 June 2019, Laxenburg, Austria.

Williams JR (1990). The erosion productivity impact calculator (EPIC) model: A case history. Philosophical Transactions of the Royal Society B 329: 421–428.

Zarei M, Izadi D & Aghababaei Samani K (2009). Detecting overlapping community structure of networks based on vertex–vertex correlations. Journal of Statistical Mechanics: Theory and Experiment. P11013.

International Institute for Applied Systems Analysis
I I A S A www.iiasa.ac.at

Address: IIASA, Schlossplatz 1, A-2361 Laxenburg, Austria

Email: rovenska@iiasa.ac.at **Department:** Communications | ERCL

Appendix A

Agenda

Workshop

Artificial Intelligence and Machine Learning for Systems Analysis of the 21st Century

24 June, 2019

C-20 Container Meeting Room, IIASA, Laxenburg

09:00-09:05	**Welcome and introduction**
09:05-11:20	**Session 1: Machine learning for classification, prediction, inference, and parameter estimation**
09:05-09:20	Miguel Poblete Cazenave (ENE) **The limits of inference without theory** Econometricians have been rightly skeptical about the usefulness of machine learning algorithms in the context of empirical work related to economic inference. Hence, I will talk about the differences in the approach to empirical work between econometrics and machine learning, the caveats of using machine learning instead of more standard approaches to econometrics, and what is the state of the art in terms of adapting machine learning algorithms for casual inference.
09:20-09:35	Steffen Fritz (ESM) **AI activities at EOCS** This presentation will give an overview of current work on ML and big data that is currently being undertaken in EOCS. In particular, it will show how big data from remote sensing is used to derive better maps of oil palm and in general improved land cover maps. The presentation also shows some work undertaken in the LandSense H2020 project as part of a quality assurance service, which automatically identifies faces and number plates to make sure citizen science data are GDPR compliant. The presentation finishes with planned work on crop type recognition from pictures using ML.
09:35-09:50	Hadi Hadi (ESM) **Mapping land cover in Indonesia at national scale using supervised machine learning algorithms and cloud computing** Hadi, Zulkarnain MT, Ekadinata A, Danylo O, Joshi N, Yowargana P, Fritz S, Kraxner F

Space to add partner logo Space to add partner logo Space to add partner logo Space to add partner logo Space to add partner logo

Degraded land map is one of fundamental data for restoration planning. This map can be derived from remote sensing data or derivative of the land cover change map. As part of the RESTORE+ project activities in Indonesia, the land cover map at national scale is needed for derived degraded land map and restoration planning at the national scale. In this contribution, we report the progress of the ongoing work to develop an operational, cloud-based, automated methodology using Google Earth Engine (GEE) platform to produce the annual, national land cover maps in Indonesia for the years 2015-2018. Supervised machine learning classification algorithms namely Random Forest and Decision Tree, which provide insights into the relative importance of the input predictors, were employed.

As input predictors, publicly-available sensor, non-sensor, and environmental data were used. Sensor data included data from optical satellites such as Landsat and Sentinel-2, as well as SAR (synthetic-aperture radar) satellites such as ALOS PALSAR annual mosaic and Sentinel-1. Using all available optical satellite images, annual composites were generated based on different criteria (e.g. median), from which spectral (e.g. vegetation indices), spatial (e.g. Haralick texture), and temporal features (e.g. harmonic model coefficients, detected change) were derived. From Sentinel-1 data, intraannual (e.g. monthly) statistical metrics (e.g. median and standard deviation) were derived. These derived predictors were used together with the original satellite data as inputs to the classification algorithm. Non-sensor data included distances to man-made infrastructure such as roads and settlements, as well as existing concession maps. Environmental data included elevation and slope. As reference data, the previous national land cover map from 2010, and additionally ground data from later years (when Sentinel satellites started to become operational) were used. A hierarchical classification scheme was employed, in which a generic set of land cover classes namely *natural forest*, *tree based system*, *non tree based system*, and *non vegetation* was further reclassified into more specific classes in several hierarchical levels. The possibilities to combine pixel-based and object-based classification approach within GEE were examined.

Preliminary (proof-of-concept) findings indicated the technical feasibility of the proposed cloud-based machine learning methodology to timely produce the updated national land cover maps, with reasonably good accuracy when the generic land cover classes were considered. However, further classification experiments regarding, among others, the input data, input data pre-processing, derived predictors, reference data, and post-processing of the predicted land cover map are needed to improve the accuracy of the produced land cover maps with detailed classes. In addition, an important next step is to perform an accuracy assessment on the predicted land cover maps, at national scale, using the established sample-based accuracy assessment framework, capitalizing on the ongoing and planned crowdsourcing activities in Indonesia to collect reference data via visual interpretation of very high resolution image data.

09:50-10:05

Andrey Krasovskii (ESM)

Application of machine learning to modeling planted forest productivity

A. Krasovskii, A. Platov, D. Schepaschenko, F. Kraxner

We will discuss the application of machine-learning methods to identify key biophysical parameters driving the forest dynamics in IIASA's global spatially explicit forestry model G4M. Practical examples for plantation forestry in Indonesia (including acacia and teak species) with special emphasis on climate impacts and soil productivity will be demonstrated.

10:05-10:20

Artem Baklanov (ASA/ESM)

Applications of machine learning to inform environmental decision making

The following results were obtained together with colleagues from IIASA and The Krasovskii Institute of Mathematics and Mechanics.

1. The first application of ML is motivated by increasing public participation in scientific research, a new global trend helping to improve existing monitoring tools. To improve the quality of data collected at crowdsourcing campaigns, we used vote aggregation procedures based on state-of-the-art machine learning algorithms and performed data pre-processing using computer vision algorithms to exclude ambiguous and low-quality images from visual inspection by volunteers. This helped significantly improve quality of the collected data.

2. The second application of ML (deep learning) is motivated by sustainability problems of oil palm expansion. Fast-growing industrial Oil Palm Plantations (OPPs) lead to significant loss of rainforest and contribute to the global warming by the corresponding decrease of carbon dioxide absorption. We proposed a novel approach to monitoring of the expansion of OPPs based on an application of a Fully Convolutional Neural Networks (FCNs) to solve Semantic Segmentation Problem for Landsat imagery. The proposed approach significantly outperforms state-of-the-art per-pixel classification methods. Moreover, the trained FCN is robust to spatial and temporal shifts of input data. This shows that FCNs enable OPPs mapping of entire countries and may serve for yearly detection of oil palm expansion.

10:20-10:35

Thomas Schinko (RISK)

Integrated assessment of short-term direct and indirect economic flood impacts including uncertainty quantification

Tobias Sieg, Thomas Schinko, Kristin Vogel, Reinhard Mechler, Bruno Merz, Heidi Kreibich

Understanding and quantifying total economic impacts of flood events is essential for flood risk management and adaptation planning. Yet, detailed estimations of joint direct and indirect flood-induced economic impacts are rare. In this study an innovative modeling procedure for the joint assessment of short-term direct and indirect economic flood impacts is introduced. The procedure is applied to 19 economic sectors in eight federal states of Germany after the flood events in 2013. Direct economic impacts are estimated by the probabilistic, tree-based Random Forests method. Our tree-based model is derived from survey data that was collected after the 2002 and 2013 flood events in Germany. The value added of this machine learning approach, compared to other existing methods for the assessment of direct economic flood impacts, is that it is object-based and considers uncertainties associated with the hazard, the exposed objects and their vulnerability. The direct economic impacts are then coupled to a supply-side InputOutput-Model to estimate the indirect economic impacts. The procedure provides distributions of direct and indirect economic impacts which capture the associated uncertainties. The distributions of the direct economic impacts in the federal states are plausible when compared to reported values. The ratio between indirect and direct economic

	impacts shows that the sectors Manufacturing, Financial and Insurance activities suffered the most from indirect economic impacts. These ratios also indicate that indirect economic impacts can be almost as high as direct economic impacts. They differ strongly between the economic sectors indicating that the application of a single factor as a proxy for the indirect impacts of all economic sectors is not appropriate.
10:35-10:50	<u>Mikhail Smilovic (WAT)</u> **A tale of two studies: Neural networks for irrigation scheduling and streamflow depletion** Federico Pasqualotto, Thomas Boerman and Mikhail Smilovic Process-based and numerical models provide the opportunity of estimating potential futures given a variety of different narratives. However, they can be both computationally intensive and necessitate providing a wide range of location-specific variables. Here we present two examples of the application of machine learning to imitate the results of such models, namely the large-scale Dynamic Global Vegetation Model LPJmL and an application of MODFLOW. LPJmL was employed to simulate wheat yields across northwestern India under ~20,000 different irrigation schedules. MODFLOW was employed to simulate the amount of streamflow depletion as the result of groundwater pumping around Lake Michigan, USA. Both studies attempt to emulate the results of the process-based and numerical models while decreasing the computational intensity and potentially allowing for predictions in areas with similar characteristics.
10:50-11:05	<u>Keivan Aghababaei Samani (EEP)</u> **Detecting community structure of complex networks by Non-negative Matrix Factorization (NMF) method** Matrix factorization methods are among common tools in Machine Learning. We use Non-negative Matrix Factorization to detect community structures in complex networks. Communities are sub-graphs of a network whose nodes are densely connected to each other, but there are much fewer connections between different sub-graphs. The method is successfully applied to a number of real-world and computer-generated networks.
11:05-11:20	<u>Christian Folberth (ESM)</u> **Combining crop modelling and machine learning for rapid provision of high-resolution crop yield predictions** Christian Folberth et al. Large-scale crop models are important tools for agricultural climate change impact assessments, as input data provider for land use modelling, and most recently for farmer information. Along these lines they can also provide information relevant for policy-making. Yet, crop modelling is highly data and computationally demanding, requiring a wide range of soil, climate, and management variables in specific formats for each target region and management system. This results in long lead times to obtain regionally or location specific information. Simultaneously, global crop simulations are frequently performed across diverse agro-environmental conditions at coarse resolutions for global impact assessments.

	To bridge the global to regional gap, we train a machine learning algorithm on global crop model outputs and input data from a 50 km resolution to produce high-resolution (1 km) crop yield estimates at the extent of Mexico as a test case. Machine learning predictions based on extreme gradient boosting compare extremely well to simulations performed at the same spatial resolution (R2=0.99, ME<1%). Essential covariates are sparser than those required by the crop model and can be provided in generic format, rendering the retrieval of crop yield predictions highly efficient. After this proof of concept, additional dimensions on crop management and an extended range of climate conditions will be included. Further exploration will cover data assimilation of crop model outputs and field observations. Crop yield predictions obtained via our machine learning approach can be operationalized for farmer information, e.g. in an app recently developed to provide location-specific yield potentials. Presently based on static simulation outputs, these can be updated on-the-fly based on farmers' input data on crop management and soil characteristics among others. Further potential applications include rapid screening of management alternatives and identification of regional hotspots for interventions in policy-making and project planning.
11:20-11:40	**Break**
11:40-12:25 *11:40-11:55*	**Session 2: Machine learning for optimization and decision support** Valeria Javalera-Rincón (ESM/ASA) **Reinforcement Learning Approach for Cooperative Control of Multi-Agent Systems** Valeria Javalera-Rincon, Vicenc Puig Cayuela, Bernardo Morcego Seix, Fernando Orduña-Cabrera Reinforcement Learning (RL) systems are trial-and-error learners. This feature altogether with delayed reward, makes RL flexible, powerful and widely accepted. However, RL could not be suitable for control of critical systems where the learning of the control actions by trial and error is not an option. In the RL literature, the use of simulated experience generated by a model is called planning. In this paper, the planningByInstruction and planningByExploration techniques are introduced, implemented and compared to coordinate, a heterogeneous multi-agent architecture for distributed Large Scale Systems (LSS). This architecture was proposed by (Javalera 2016). The models used in this approach are part of a distributed architecture of agents. These models are used to simulate the behaviour of the system when some coordinated actions are applied. This experience is learned by the so-called, LINKER agents, during an off-line training. An exploitation algorithm is used online, to coordinate and optimize the value of overlapping control variables of the agents in the distributed architecture in a cooperative way. This paper also presents a technique that offers a solution to the problem of the number of learning steps required to converge toward an optimal (or can be sub-optimal) policy for distributed control systems. An example is used to illustrate the proposed approach, showing exciting and promising results regarding the applicability to real systems.

11:55-12:10	Tatiana Ermolieva (ESM) **Artificial Intelligence: non-smooth stochastic optimization and iterative quasigradient procedures for machine learning** Tatiana Ermolieva, Yuri Ermoliev, Petr Havlik, Esther Boere, Juraj Balkovic, Michael Obersteiner The talk illustrates the importance of new type non-smooth stochastic optimization and stochastic quasigradient (SQG) procedures for robust on-line and of-line (real-time operational and strategic) decisions involving large-scale machine learning and Artificial Intelligence problems in particular, deep learning including deep neural networks learning. Proposed models and methods enable to use probabilistic and non-probabilistic (explicitly given or simulated) probability distributions of stochastic parameters/data, combining measures of chances, experts believes and similarity measures (for example, compressed form of the kernel estimators), as well as non-smooth and discontinuous regret/reward/goal functions. This is vitally important for integrative sustainable developments modeling. We highlight the need for these approaches using case studies on food-water-energy-environmental security NEXUS management characterized by complex, often unknown and unobservable explicitly, interdependencies between systems and regions at different resolutions, in the presence of risks and uncertainty, uncertainty about uncertainty, incomplete and asymmetric information.
12:10-12:25	Finn Laurien (RISK) **Making smarter and faster decisions based on machine learning** Finn Laurien, Stefan Velev, Stefan Hochrainer-Stigler In the digital age, where buzzwords such as artificial intelligence and machine learning are omnipresent, social scientists are confronted with increasing challenges of monitoring and assessing larger and more complex sets of data collected by qualitative or semiquantitative methods. Semantic text analysis methods based on machine learning algorithms can guide social scientists to better understand, structure, cluster and assess the large amount of data. We apply a semantic text analysis method with more than 30.000 user feedbacks and comments to assess the user perception and the usability of the Flood resilience measurement tool. The approach is able to highlight strengths and weaknesses for specific measures collecting complex system elements of community flood resilience. Based on the results we argue that semantic text analysis methods should be implemented more often to help prioritizing qualitative results but also to aid trust and validity to quantitative methods to validate their outcomes for policymakers.
12:25-13:00	**Discussion**
13:00	**Close**