Report from Dagstuhl Seminar 22091

AI for the Social Good

Claudia Clopath[*1], **Ruben De Winne**[*2], **and Tom Schaul**[*3]

1 **Imperial College London, GB.** c.clopath@imperial.ac.uk
2 **Oxfam Novib – The Hague, NL.** ruben.dewinne@oxfamnovib.nl
3 **Google DeepMind – London, GB.** schaul@google.com

Abstract

Progress in the field of Artificial intelligence (AI) and machine learning (ML) has not slowed down in recent years. Long-standing challenges like Go have fallen and the technology has entered daily use via the vision, speech or translation capabilities in billions of smartphones. The pace of research progress shows no signs of slowing down, and demand for talent is unprecedented. AI for Social Good in general is trying to ensure that the social good does not become an afterthought, but that society benefits as a whole. In this Dagstuhl Seminar, which can be considered a follow-up edition of Dagstuhl Seminar 19082, we brought together AI and machine learning researchers with non-governmental organisations (NGOs), as they already pursue a social good goal, have rich domain knowledge, and vast networks with (non-)governmental actors in developing countries. Such collaborations benefit both sides: on the one hand, the new techniques can help with prediction, data analysis, modelling, or decision making. On the other hand, the NGOs' domains contain many non-standard conditions, like missing data, side-effects, or multiple competing objectives, all of which are fascinating research challenges in themselves. And of course, publication impact is substantially enhanced when a method has real-world impact. In this seminar, researchers and practitioners from diverse areas of machine learning joined stakeholders from a range of NGOs to spend a week together. We first pursued an improved understanding of each side's challenges and established a common language, via presentations and discussion groups. Building on this foundation, we organised a hackathon around some existing technical questions within the NGOs to scope the applicability of AI methods and seed collaborations. Finally, we defined guidelines and next steps for future AI for Social Good initiatives.

Seminar February 27–4, 2022 – http://www.dagstuhl.de/22091
2012 ACM Subject Classification Computing methodologies → Machine learning
Keywords and phrases Machine Learning, Artificial Intelligence, Social Good, NGO, sustainable development goals, Non-governmental organisation
Digital Object Identifier 10.4230/DagRep.12.2.134

1 Executive Summary

Ruben De Winne (Oxfam Novib – The Hague, NL, ruben.dewinne@oxfamnovib.nl)

AI and ML have made impressive progress in the last few years. Long-standing challenges like Go have fallen and the technology has entered daily use via the vision, speech or translation capabilities in billions of smartphones. The pace of research progress shows no signs of slowing down, and demand for talent is unprecedented. But as part of a wider AI for Social Good trend, this seminar wanted to contribute to ensuring that the social good does not become an afterthought in the rapid AI and ML evolution, but that society benefits as a whole. The

[*] Editor / Organizer

five-day seminar brought together AI and ML researchers from various universities with representatives from NGOs based in Benin, Tanzania, Uganda, The Netherlands and globally. These NGOs all pursue various social good goals, such as improving air quality, increasing agricultural productivity with the help of technology, transforming health care, providing humanitarian support, and defeating poverty. On these topics, NGOs have rich domain knowledge, just like they have vast networks with (non-)governmental actors in developing countries. Mostly, NGOs have their finger on the pulse of the challenges that the world and especially its most vulnerable inhabitants are facing today, and will be facing tomorrow. The objective of the seminar was to look at these challenges through an AI and ML lens, to explore if and how these technologies could help NGOs to address these challenges. The motivation was also that collaborations between AI and ML researchers and NGOs could benefit both sides: on the one hand, the new techniques can help with prediction, data analysis, modelling, or decision making. On the other hand, the NGOs' domains contain many non-standard conditions, like missing data, side-effects, or multiple competing objectives, all of which are fascinating research challenges in themselves. And of course, publication impact is substantially enhanced when a method has real-world impact. The seminar facilitated the exploration of possible collaborations between AI and ML researchers and NGOs through a two-pronged approach. This approach combined high-level talks and discussions on the one hand with a hands-on hackathon on the other hand. High-level talks and discussions focused first on the central concepts and theories in AI and ML and in the NGOs' development work, before diving into specific issues such as generalisability, data pipelines, and explainability. These talks and discussions allowed all participants – in a very short time-frame – to reach a sufficient level of understanding of each other's work. This understanding was the basis to then start investigating jointly through a hackathon how AI and ML could help addressing the real-world challenges presented by the NGOs. At the start of the hackathon, an open marketplace-like setting allowed AI and ML researchers and NGOs to find the best match between technological supply and demand. When teams of researchers and NGOs were established, their initial objective was not to start coding, but to define objectives, assess scope and feasibility. The intense exchanges during the hackathon allowed NGOs with a lower AI/ML maturity increased to increase understanding of the capabilities of AI/ML and define actions to effectively start working with AI/ML. NGOs that already had a more advanced understanding and use of AI/ML technology prior to the seminar, could take their AI maturity to the next level by trying out new ML approaches, designing and testing tailored ML models, or simply exploring new partnerships. Key to this success of the hackathon – and the seminar at large – was the presence of AI/ML experts whose respective fields of expertise could seamlessly be matched with the various needs of the various NGOs. This excellent group composition also facilitated a productive discussion about guidelines on how to do effective AI for social good collaborations in the future (e.g. by focusing on long-term partnerships, and by sequencing problem scoping before data cleaning and – only in last instance – an actual hackathon).

2 Table of Contents

3 Overview of Talks

3.1 Introduction to the seminar

Claudia Clopath (Imperial College London, GB)

Claudia Clopath introduced the notion of "AI for Social Good" and the motivation for the seminar. In particular, she listed the following goals: 1) Bring NGO participants and AI researchers together, 2) Create awareness, 3) Make sure incentives are aligned, 4) New innovations coming out of the newly formed collaborations. The potential outcomes of the seminars were: 1) Understanding the link between the two worlds: Machine Learning and NGOs, 2) Understanding the ways to make AI for Social Good initiatives work, 3) Build Prototypes during the hackathon time, 4) Build collaborations, create a network, make friends, 5) and prepare a set of follow-up plans. Finally, she defined a set of rules for the seminar: Connect, Respect, Bottom-Up, Brainstorm, Fun.

3.2 Introduction to Artificial Intelligence

Tom Schaul (Google DeepMind – London, GB)

Tom Schaul gave a brief overview of Artificial Intelligence. He talked about the promise of AI, gave an overview of the subfields of AI terminology, explained some key concepts specifically around the types of machine learning (i.e. supervised learning, unsupervised learning, and reinforcement learning), as well as the most common issues (grouped into data, model and process challenges). An important comment made during the session was that a human needs to check results when deploying machine learning in the real world. In addition, NGO participants in the session were advised that simple but trustable, interpretable and robust methods should always be the first they start with. Only when simple methods aren't sufficient, then one can go to more complex methods such as neural networks. Participants were also warned about the risk of overfitting, which happens when the model becomes so good, doesn't make any mistakes anymore, and gets too precise, and gets bad at anything else. E.g. when detecting a blue pixel in images with cats leads to predicting a "cat" on images with a blue pixel but no actual cat. To mitigate the risk of overfitting, the choice of a validation set is very important: it is advised to use a validation set that uses data from the real-world use case the NGO is working with.

3.3 Deploying computer vision for conservation

Sara Beery (California Institute of Technology – Pasadena, US)

Sara Beery presented a use case of AI for social good, i.e. the deployment of computer vision for ecological conservation. She addressed challenges in deploying impactful computer vision solutions for problems that are faced by conservation land managers, noting the importance of also being honest about failures. She pointed to usefulness, accessibility, collaborative development and feedback, and strong communication of risks and intended use as the main factors that enable impact of computer vision in the field of ecological conservation. In addition, interdisciplinary collaboration was mentioned as a factor leading to deeper understanding and new interesting challenges for computer vision. To conclude, Sara emphasized the need for a sustainable plan for supporting any tools that were developed. Indeed, long-term support is needed, as a computer vision model in itself is not a solution. For NGOs, it is critical to empower them by building capacity within the community and/or by advocating for the community to hire the right capacity.

3.4 Dark ecology: unraveling mysteries of bird migration using RADAR and Machine Learning

Subhransu Maji (University of Massachusetts – Amherst, US)

Subhransu Maji started his presentation by explaining that billions of birds migrate every year, mostly in the cover of darkness. But these birds are visible on RADAR networks in the continental US. Subhransu explained that using Machine Learning, we can learn how migration has changed over the last 25+ years. A team of ecologists and computer scientists worked together to analyze this bird migration data at scale. Challenges and unique opportunities that this collaboration had were also discussed. Subhransu concluded with three pieces of advice to the participants:

- Don't throw away noisy data! You might be able to correct for noise.
- Don't throw away info on who labeled data!
- Don't throw away intermediate things, might be useful at some point for training.

4 Working groups

The seminar agenda left ample room for spontaneous discussions. A brainstorm session led participants to select the following topics for further exploration in smaller working groups: Off-the-shelf AI apps, tools, and platforms, and what holds NGOs back from using these?

- Task-generalizability and regional generalizability
- Setting up data pipelines connected to model deployment, i.e. getting data, labeling data, how to get data pipelines up and running
- What methods exist to understand text
- Tools to support explainability and decision-making
- ML project management: feedback, sustainability, incentives, defining the problem

At the end of these discussions, the groups came back together in the plenary session and shared their main conclusions with each other. The group that had discussed generalizability for instance, noted that generalizability of your model should never be assumed, but should always be tested. The group that discussed data pipelines fed back to the wider group that certain data sets are made available, e.g. by governments/government-funded agencies, and that dealing with anomalies is important, which is especially true for sensor data (because sensors sometimes break etc.).

5 Hackathon

During the hackathon, which was spread out over two days, five groups consisting of machine learning experts and NGOs tackled the real-world issues that the latter had brought to Dagstuhl. As several NGOs had brought issues that required similar technological expertise, they joined forces and benefited from each other's perspectives, challenges, and lessons learned.

Group 1 AirQo from Uganda and Laterite from the Netherlands brought seemingly different cases to the table, i.e. measuring air quality more efficiently with a limited number of sensors placed at various locations across Uganda (fixed locations as well as in moving vehicles), and predicting school drop-out on the basis of various sources of data, for example, surveys of populations, respectively. Nonetheless, in machine learning terms what unified both cases was that their data was feature-based (and not, for example, image- or text-based). Machine Learning algorithms such as Gaussian Processes (GPs) and Gradient Boosted Decision Trees (GBDTs) perform well for such use-cases. Before actually applying these algorithms to their respective cases, both AirQo and Laterite needed to rephrase their problem and properly set up their data pipeline. Laterite went with using GBDTs while AirQo used GPs to balance various trade-offs of these algorithms. At the end of the hackathon, different approaches had been tested with real-world data, and the group had agreed to continue their collaboration beyond the seminar.

Group 2 Soon dubbed "the text group" – Oxfam Novib, Save the Children, and the Red Cross joined forces to discuss and address challenges with natural language. While Oxfam Novib and Save the Children aimed at automating knowledge management and reporting, the Red Cross wanted an algorithm to classify open-text survey responses. Oxfam Novib and Save the Children invested their time mostly in scoping their problem, conducting exploratory online interviews with internal stakeholders, and designing a plan for action after the seminar. The Red Cross sub-team did develop a prototype algorithm that could effectively do the desired classification. The team also considered applicability of the model beyond the particular domain it looked at for this case (i.e. rumors and opinions about COVID-19). It documented the results of tests with various models. And last but not least, it secured internal funding at the Red Cross to continue working on this project.

Group 3 TechnoServe, Humanitarian OpenStreetMap, and again the Red Cross all brought cases requiring computer vision technology. This group probably got furthest with developing and testing actual machine learning models. At the end of the hackathon, their model could indeed recognize trees (catering to TechnoServe's case) and recognize buildings (useful for the cases of Humanitarian OpenStreetMap and the Red Cross). The machine learning experts in the group committed to remaining available for questions after the seminar, and the NGOs would give a demonstration of the models in their

respective organizations.

Group 4 In the fourth and last group, D-tree from Tanzania, which has strong in-house machine learning expertise, mostly scoped out ideas of how to further advance their use of machine learning. Together with the machine learning experts at the seminar, they did two iterations on these ideas with the team in Tanzania. This resulted in a focused list of ideas that D-tree could pick up after the seminar (e.g. solving the interpretability problem, tailoring their health survey, designing a more continuous instead of binary scale, personalizing incentives, detecting suspicious visits, and predicting the type of intervention instead of predicting the outcome). For most of these ideas, simple solutions such as software engineering could already be very valuable. To bring these ideas into practice, the group also identified possible collaborations and partnerships in this area.

All in all, the hackathon was a success: NGOs with a lower AI/ML maturity increased their understanding of the capabilities of AI/ML, while NGOs that already had a more advanced understanding and use of AI/ML technology could take a next step. Key to this success was the presence of AI/ML experts whose respective fields of expertise could seamlessly be matched with the various needs of the various NGOs. In times of COVID-19 – with the many impediments to international travel and even a fair share last-minute cancellations – the seminar was fortunate to have this nearly perfect match between offer and demand of skills and expertise.

6 Guidelines on how to do effective AI for social good collaborations in the future

On the final day of the seminar, the participants reflected on the success of the seminar and on what had been the enabling factors. They formulated a set of guidelines on how do effective AI for social good collaborations in the future. For seminars like the AI for Social Good seminar at Dagstuhl, the importance of in-person attendance was underscored. In addition to AI/ML researchers and domain experts, software engineers should be invited, and the presence of financial partners could be considered. Also, the seminar programme could include a session on what AI can and cannot do, if several NGO participants lack this knowledge. For AI/ML participants, a session or a talk at the start of the seminar about structures and constraints for NGOs could be useful. In particular, domain experts should "educate" AI experts on what is needed/feasible in the field, so that expectations of both worlds can be better aligned. Concretely: AI experts should share experiences from previous (failed) AI for social good pilots/experiments/projects. For visibility purposes, it's advised to broadly advertise AI for Social Good seminars to NGOs (e.g. through NetHope, or at individual organizations' events) and to invest in communication via social media (e.g. by publishing success stories of seminar 22091 to justify NGOs' time). In general, AI for social good collaborations should be clear from the outset about the willingness on both sides to pursue the collaboration in the long-term. Long-term partnerships should be established between affiliations/organizations, not (only) between individuals. But attention point: partnerships between organizations can be painful to set up. To facilitate long-term partnerships, a GitHub team page/repository could be set up with a to-do list per project, so that others can pick up work when they have time, or at another workshop etc. A concrete avenue here would be to tap into existing summer schools or other programmes of universities related to AI for social good (e.g. UMass Amherst, University of Washington, DSSG, Data

summer school at Ecole polytechnique in France, DeepLearning Indaba). More long-term partnerships with academics can also be very useful, because they have lots of students who need lots of projects. The challenge here is that the students are junior, so the senior academic needs to buy into the collaboration and supervise their students.

AI for social good collaborations should ideally follow the following sequence: 1) a scoping exercise to clearly define the problem and identify potentially relevant datasets; 2) a "cleanathon" where software engineers look at the data, 3) the actual hackathon(s) where algorithms/models are designed and tested. In-person interactions are deemed effective for all three of these phases.

Finally, it was suggested to bring NGOs to AI conferences and workshops (e.g. the likes of NeurIPS, etc.).

7 Follow-up actions

To close the seminar while opening concrete perspectives for longer term collaboration, the four groups agreed on action plans with the following structure:

- What to do next week
- What to do next month
- What to do next year
- Who does what?

The entire group also agreed to explore collaboration with HuggingFace, which could host AI for social good models and provide additional (paid) services. They also agreed to have an (online) follow-up session in September 2022, during which they would honestly share experiences with each other and pitch the results achieved at the Dagstuhl Seminar to funders.

Participants

- Emran Alchikh Alnajar
Humanitarian OpenStreetMap, DE
- Asma Atamna
Ruhr-Universität Bochum, DE
- Sara Beery
California Institute of Technology – Pasadena, US
- Claudia Clopath
Imperial College London, GB
- Ruben De Winne
Oxfam Novib – The Hague, NL
- Jose Pablo Folch
Imperial College London, GB
- Jörg Franke
Universität Freiburg, DE
- Mariella Goebl
TechnoServe – Cotonou, BJ
- Nitusima Kataraia
D-tree – Dar es Salaam, TZ
- Stijn Koster
Save the Children – The Hague, NL
- Rik Linssen
Laterite – Amsterdam, NL
- Subhransu Maji
University of Massachusetts – Amherst, US
- Jacopo Margutti
510 / Netherlands Red Cross – The Hague, NL
- Corina Markodimitraki
510 / Netherlands Red Cross – The Hague, NL
- Nele Quast
University of Oxford, GB
- Raghu Rajan
Universität Freiburg, DE
- José Ricardo Rubio Valverde
Laterite – Amsterdam, NL
- Samira Said
D-tree – Dar es Salaam, TZ
- Tom Schaul
Google DeepMind – London, GB
- Ruby Sedgwick
Imperial College London, GB
- Richard Sserunjogi
AirQo – Kampala, UG
- Christoph Weigl
TechnoServe – Arlington, US