



Designing Data-Driven Solutions to Societal Problems: Challenges and Approaches

Janardan Misra^(✉), Divya Rawat, and Sanjay Podder

Accenture Labs, Bangalore, India
{janardan.misra, d. b. rawat,
sanjay.podder}@accenture.com

Abstract. Deciding effective and timely preventive measures to address societal problems is a difficult challenge. Societal problems tend to be inherently complex, affect the bottom of the socio-economic pyramid, and have wider spread across populations and geographies with resource constraints. In this paper we identify challenges and propose certain design considerations to guide development of machine learning based data-driven applications for addressing societal problems. Challenges cover subtle difficulties, which may be encountered during various phases of design and deployment life-cycles and design considerations make recommendations with respect to currently known state-of-the-art concepts, tools, and techniques to address these challenges.

Keywords: Societal problems · Data-driven solutions · Machine learning

1 Introduction

Innovations in digital technology have steadily transformed various information centric processes across the globe affecting almost all walks of life. However, effectiveness of existing digital solutions found in governments and other social service organizations is being increasingly questioned considering the observations that when an event takes place, such solutions often do not enable efficient coordination and collaboration among various entities of interest and are not designed to solve complex scenarios which may arise in practice [1]. Primary reason for such inefficiencies is that these solutions are based upon slightly outdated technologies and design thinking and do not take advantage of recent innovations in the information and communication technology (ICT) space including cloud computing, mobile connectivity, and analytics driven decision making.

At the same time, relatively low-income societies still face a multitude of challenges including low empowerment of weaker sections of society, poor health and low nutrition, low quality of education, poor child protection, and poor quality of sanitation and hygiene [2]. To address these challenges and resulting societal problems like child trafficking [3, 4], there is a need to invent novel solutions applying ICT involving social, mobile, analytics, and cloud based digital technologies.

Recent studies [5, 6] discuss design approaches for implementing digital solutions for delivering high-quality outreach services. For example, to overcome the challenge of relatively lower internet penetration in rural areas, organizations in social service

sectors, have started adopting mobile based decision support systems (MDSS) that can work without requiring internet connectivity. MDSS have helped many organizations catering to outreach care, with in-built rule sets to categorize the target population and ease the work of outreach workers from complex analysis based upon multiple guidelines. Thus, penetration of low cost mobile devices in rural areas has started enabling organizations in social-sector to upskill the outreach workers through digital technologies.

Continuing with [7], in this paper, we focus on data-driven, machine-learning (ML) based, dynamic and context aware computational models, which could help to provide improved quality of solutions especially for relatively complex scenarios. For example, in [20] it was reported that semantic analysis of large number of mobile based text messages sent by teens to volunteers revealed clear patterns, which could be used in assisting volunteers to judge criticality of situation in which a teen might be and respond with higher effectiveness to resolve the problem.

We discuss challenges which ML designers may encounter during various phases of design and deployment life-cycles of the applications aimed for helping those working on addressing societal challenges and make recommendations with respect to currently known state-of-the-art concepts, tools, and techniques in ML to address these challenges. For example, a key challenge while designing data-driven computational solutions in the context of social problems is the lack of verifiable and quality data. This is because owing to various socio-economic constraints, researchers often rely on non-governmental organizations (NGOs) as a primary source of data, however the data collected by NGOs may not be well suited for ML based applications designed to extracting useful patterns from the data.

The paper is organized as follows: Sect. 2 presents discussion on difficulty of solving societal problems to design social good applications and continues with challenges which ML designers may encounter in practice. Next, in Sect. 3, we present high-level design considerations, which may help taking design decisions at various stages. Section 4 presents concluding thoughts.

2 Challenges in Designing ML Models for Social Problems

2.1 Inherent Hardness of Modeling Social Phenomena

To design solutions for societal challenges using traditional approaches of building digital applications, it is essential to model underlying problems and propose solutions to those problems analytically with help from field experts or social scientists. An example of this was presented in [9], wherein for the problem of identifying childhood vulnerabilities including trafficking, authors proposed a linear convex model with 32 features along with a threshold to determine whether a child is vulnerable or not.

However, social phenomena are inherently hard to model accurately [5, 8]. The primary reason for this could be attributed to large number and variety of factors affecting the phenomena under study in ways too complex to be fully understood. To further complicate the matter in the context of social problems, for ethical reasons, controlled experiments cannot be performed since actual negative social events cannot

be artificially created but could only be analyzed when they occur naturally. Therefore, solutions based upon manual analytical approaches (e.g., closed form formula-based vulnerability analysis [9]) cannot reliably generalize to larger contexts and might remain locally relevant where most of the parameters in the model are approximately fixed and attributes with high predictive power are known with field experience.

When generalization beyond local sociocultural boundaries and large-scale adoption are critical goals to achieve, a data-driven machine learning based approach may provide an effective work around to this problem. Under such design framework, a computational model is generated (instead of a manually defined analytical model) from sample data collected from the field studies with feature-set designed in consultation with social scientists specializing in that field.

Towards that, we aim to evolve an approach towards building a design-framework for applications aiming to address wide spectrum of social problems especially affecting bottom of the socio-economic pyramid and having wider spread across populations and geographies with resource constraints. Primary objective is to apply data-driven design methodology together with application of ML techniques to render eventual solution amenable to wider adoption with low cost imprint and serving priorities at multiple levels ranging from potential victims (e.g., children as potential targets of trafficking) to field workers, to NGOs and Government Agencies interested in analysis of impact of their services, and eventually to social scientist interested in scientifically studying the underlying phenomena at larger scales.

In the following, we will refer these applications as DDSSP (data-driven solutions for societal problems).

2.2 Challenge: Issues with Data

One of the major problems in designing DDSSP applications lies in issues related to availability of data suitable for training a ML tool. Drawing from our experience, data quality issues may arise from following scenarios:

1. *Non-standardized labels for dependent variables:* This is a big problem as having a labeled data to start with is important to train supervised ML models. Even when a questionnaire in a survey is pre-defined, the fields which may be of use for ML design may depend on the perceptions and writing style of the surveyors (for example suggested actions). This adds to the task of standardizing the labels based upon the understanding of target social problem for which ML model is to be build.
2. *Lack of structured format:* Data related to social issues is often collected in form of field surveys. However, these surveys may not have been designed for their eventual use in building data-driven analytics applications. For example, the information collected over surveys might be in arbitrary form – it may contain highly verbose descriptions for the categorical variables. Else details may be present in such a form that without manually understanding contents, it is difficult to organize them properly.
3. *Missing criteria for unique Identification of data points:* It is possible that there are no fields to identify data points uniquely. For example, using names to refer to people may work in surveys collected at local level by a field agent, however, when

data sets from different field agents are combined, names may not be able to uniquely differentiate all.

4. *Missing time-stamps*: Survey data collected over time need to be time stamped and should be stored in a way that temporal changes over different time points are inferable. However, in practice, such time stamps might be missing, incomplete, or implicitly recorded (for example, as a part of file names “Survey_JM_12Dec.docx”).

With all the above stated problems and lack of data governance standards, data collected may not be readily suitable for ML applications to design and work on. For example, following issues were observed in the data sets we encountered: Data was dispersed in an unorganized manner over multiple excel files and SQL dump. There was no meta-data about the SQL database to relate or map the information across tables. Also, unique identifier or primary key was missing in most of Excel and SQL data. Only after manually going over the tables in the dump and experimenting with different combination of fields in the table as keys, we could merge tables only partially incurring significant loss of data.

Sub-challenge: Semantic Inconsistencies. It is not uncommon to have inconsistencies in target variables selected for ML design from the data collected from field-agents. Such inconsistencies can reduce accuracy of the ML solution to low levels and therefore require additional approaches for their correction. One approach is to build interpretable ML models using original data and manually inspect the learned ML model for the consistency. For example, extracting classification rules from the Random Forrest classifier and analyzing all related rules together.

Sub-challenge: How to Use ML Outputs in Practice? When addressing social problems, the cost of false negatives (e.g., girls which are incorrectly declared as not vulnerable to trafficking but actually are) is high and in some sensitive cases is not even compensable. Therefore, role of field agents becomes essential and ML models should primarily be deployed to assist field agents in making informed decision instead of taking actions based upon the recommendations given by ML based solution.

Sub-challenge: How to Deal with Missing Labels? Cases of missing labels in data are generally ignored from the training set during ML model designs but in DDSSP applications it is recommended that these cases are treated as exceptional cases in which it might have been difficult for a field agent to take decision and where explicit expert intervention may be needed and therefore such examples could be classified under a new class “Expert Intervention Needed”.

2.3 Challenge: Missing Semantic Clues from the Context

Once the data is prepared for training a ML model, often designers are left with small number of data points to start with. Effective learning from small amount of data is difficult as well as prone to incur biases (i.e., low generalizability). With limited datasets, it also becomes difficult to decide on the ML technique to use. On the other hand, generating more data synthetically is difficult as it would require modeling accurately underlying socioeconomic and cultural factors. Coupled with the problem of sparsity of data, is the problem of absence of details from the temporally relevant context in which

social problems occur. These missing details may sometime contain actual causal factors, which might be contributing to the occurrence of the problem. One such example is the occurrence of recent disturbances in the family, which might have mentally agonized a child and in turn made him/her vulnerable to the trap of anti-social elements involved in trafficking.

In contrast, human analysts can learn a great deal about the underlying problems even with just few instances alone because they can associate information contained within these examples with the semantic context in which these events occurred and using their latent expertise on the subject matter and commonsense reasoning, they can arrive at correct remediation strategies.

Therefore, it might help to explore ML techniques which are known to deal better with small data sets (see [24–26]) or devise new one which can learn with very few data points to start-with. Also, advancements in the field of commonsense based ML reasoning [28, 29] should help bring positive value in the design of applications in social domain.

Furthermore, prior identification of measurable semantic signals from the environment, which might be playing subtle roles in a social problem, with the help from those deeply involved in the actual field work, would help in designing ML models with high accuracy. Feature engineering is known to be play central role in ML and for social problems it appears critical.

2.4 Challenge: Geographical Differences Do Matter i.e., Designing ML Models by Combining Data from Different Sociocultural Regions or Demographics May not Yield Reliable Solutions

As the demographics largely play an important role in defining the nature and causes of a social problem; the data collected over different sources/regions mostly results in having variations in core-elements (like causes, effects etc.) of the same social problem. Therefore, there is high possibility that ML model when trained on data from one region will not perform as expected when given data arriving from a different region.

For example: A region X with no educational institution in vicinity may have contributed to child labor, which in turn might have made children vulnerable for trafficking, whereas in a different region Y, cultural biases against girl children might be rendering them vulnerable for trafficking. Therefore, if a ML model is trained to predict vulnerability of a child using data from region X, it may fail to generalize well when applied on region Y (and vice versa).

2.5 Challenge: Difficult to Learn Complex Relationship of Events, Human Behavior, and Decision Making

Identifying social issues in human societies is an inherently hard problem because of the existence of complex relationships among various elements in the society. The social issues could be a result of different independent choices taken by the individuals or groups or whole institutions over a span of time.

For example, a girl being vulnerable to certain kind of issue largely depends on her relationship with the surrounding elements like being the eldest among the siblings she

might have to stay back at home to look after them while her parents go for work or in a different setting a girl might have to go to some low wage job to help the family financially.

Enabling automated learning of such subtle factors and complex relationships, building over a span of time to is a difficult objective for state of the art ML systems. Lifelong learning is an emerging area which might bring potential to address such challenges [27] if ML application is getting deployed very close to potentially vulnerable populations.

2.6 Challenge: Difficulty of Carrying Out Experimental Pilot Studies

Collecting details about actual victims of social problems is a known challenge [5] – primarily because these victims are generally out of access for detailed examination and only indirect data points could be collected with enough efforts. On the other hand, data for non-victims is relatively easier to acquire but it only makes design of prediction model harder owing to inherent bias towards non-victim class. Additional difficulty arises because when a prediction model is used in practice, its predictions control mitigation strategies which further biases population towards its predictions and hence make it harder to know to what extent such a model is inherently accurate.

Conducting pilot studies to estimate and improve on the efficiency of ML models is very difficult as it requires mimicking a real-world scenario with respect to social problem. This could be impossible in certain cases like human trafficking. Alternately, estimating ML model's performance on real scenarios would require carrying out extensive surveys over a span of time including cases where potentially vulnerable ones actually became victims, which is an inherently difficult process requiring extensive support – something not easy to find.

2.7 Challenge: Offline Models Versus Explicit Representations of Models

As low-income geographies with high resource constraints have relatively higher incidents of the social problems, it may not help to build ICT solutions which require heavy computational machinery or networking support for deployment. These sections of societies might be deprived of even basic networking facilities. For example, as per the World Energy Outlook [12], as of 2016, 33% of rural areas in developing countries had no electricity. In such cases offline light weight pre-trained models are more useful for actual usage.

Another alternative is to have an explicit representation of trained ML Model, which can be embedded in simpler forms into the mobile application. For example, extracting logistic regression equations from the trained model and using these equations directly to compute the confidence scores for the target potential vulnerabilities and recommended mitigation programs when data for a new case is encountered during actual usage.

3 Design Considerations for ML Based Data-Driven Applications

3.1 Design Consideration: Predicting Vulnerabilities from Data Eventually

Designing solutions for complex social problems with detailed manual analysis is inherently hard and error prone. An effective alternative is to design a model which optimally conforms to the data collected from the real scenarios. Machine learning based techniques provide operational solution wherein patterns underlying the data related to actual instances of the problems could provide clues to solving the problems computationally and in designing mitigation strategies.

- *Machine learning based predictive modelling for deciding preventive measures:* Often solving social problems requires an ability to make predictions well ahead of time before actual negative event may take place (e.g., vulnerability prediction for child trafficking problem) using analysis of factors affecting potential victims. In this perspective classification and regression techniques may be used to design required predictive model though initial design trials may be necessary to determine the right prediction technique or a combination of many [10].
- *Dealing with Cold-Start problem:* However, acquiring sufficient good quality data to train machine learning models in the context of social problems is difficult. This may result into *cold-start* problem if only ML based model must be used to design DDSSP applications. For this reason, ML based data-driven solution should be the *eventual* design goal and in-order to start its application in the field work, one needs to have alternative solutions resulting from prior field experiences designed in collaboration with social experts.

3.2 Design Consideration: Use Structural Patterns in Data for Planning Actions

Similarity Analysis. Similarities among potential victims can be used to identify social-groups and to identify outliers. For example, a *critical-vulnerability profile* (CVP) containing only those factors which may render a potential victim highly vulnerable could be defined and all the known respondents having similar CVPs within same locality can be made to socially connect with each other so that they can work as a group to address their vulnerabilities together.

Clustering Analysis. Clustering analysis can be used to determine whether certain details about a new respondent are far away from others in the same locality? Note that in low income geographies, high levels of social similarities within same locality are a commonly observed phenomenon. If so, DDSSP application alerts the agent with factors where high deviations are present.

Contextual Modelling. The similarity graphs or clusters can be further augmented with contextual knowledge about external environmental factors affecting the underlying phenomena (e.g., large scale religious gathering making trafficking of children

easier for anti-social elements [4]). Such augmented graphs (type of knowledge graphs) can further assist in taking timely preventive measures as per the emerging contexts.

3.3 Design Consideration: (Causal Inference) Make Decisions only Based upon the Causal Analysis of the Effectiveness of Actions in Past

After identifying potential vulnerabilities using predictive modelling, next logical step is to determine mitigation strategies to reduce existing vulnerabilities of potential victims. Here again, data driven statistically sound approaches should be applied to first estimate relative effectiveness of different mitigation programs and based upon that provide recommendations.

However, establishing causal associations between mitigation programs and reduction in vulnerabilities is difficult since it would require careful analysis of statistically significant amount of data from *randomized trials* [13, 14] involving the cases where a mitigation program was enforced (treated group) and where no mitigation program was enforced (control group). Identification of *confounding variables* to explain actual outcomes is yet another challenge to deal with in such analysis which would require adoption of methods like multivariate modelling or propensity scores [15] together with an intervention of subject matter experts.

Even though such rigorous studies may be time consuming as well as expensive, from the perspective of long term, large scale impact, they need to be given due consideration. As an illustrative example, authors in [16] identify those villages and poor households who genuinely require help such that given help would do something positive than what they could have done themselves. Towards that they conducted a randomized controlled trial at village and household levels to study effect of unconditional cash transfers on psychological well-being and food-security.

In situations where data from randomized experiments is not available and only observational data is accessible, techniques [17, 18] like Additive Noise Methods (ANM), Information Geometric Causal Inference (IGCI), difference-in-differences (DID) analyses, instrumental variables (IV), and Regression Discontinuity Designs (RDD) can be used in conjunction with tools like causalImpact [19].

3.4 Design Consideration: (Continuous Learning) Design Applications with Components for Continuous Learning Based Dynamic Evolution of ML Models

To motivate this design choice, let us consider a hypothetical scenario related to human trafficking use case. In this scenario, let us assume that there has recently been cases of child trafficking in a locality during a large gathering, however, not all of those victims were correctly predicted to be vulnerable by the existing model. Therefore, to update underlying prediction model, new data needs to be sent to its designers, which would then involve new cycle of update and reloading of the predictive model to the agent devices on periodic basis. Often such solutions even if built using ML techniques require centralized offline update of the predictive model and DDSSP applications running on agent devices cannot adapt themselves at run-time when new cases of actual victims become known!

Towards that we suggest that solutions for social problems must be designed as continuously adaptive applications which learn (from potentially incomplete data) while being in actual use by retraining themselves automatically when information about new actual incidents is entered on the agent device running the application. Eventually overtime each agent would have evolved its own unique predictive model based upon the incidents of the trafficking known in her area and other cases where such trafficking did not take place for known period. Applications should also update their prior predictions after improved training and send alerts about all those, who now are in danger zone but earlier were not.

Additionally, agent device or central server should be designed to analyze updated field-data to infer which factors are becoming increasingly critical in the light of new incidents so that right mitigation strategies can be designed or existing ones could be adapted to meet the requirements of the emerging scenarios. For example, based upon these updated predictions, DDSSP application for child trafficking should send alerts to all the registered children (and/or their care takers) and community facilitators regarding changes in the mitigation strategies.

As discussed before, Lifelong learning [27] is an emerging area which might bring potential to offer such features if ML application is designed such a way that it can accept details from field agents in natural conversational forms and then automatically extract relevant data for self-retraining.

3.5 Design Consideration: (Network Effect) Include Features to Enable Network Effect for Collective Collaboration

ICT should be effectively used to connect various human elements (e.g., potential victims, CFs, governance bodies) and computing devices with each other on larger scales across regions and communities in order to collectively unite and work against the root causes of the problems in a way which is more effective than what could have been achieved without being connected at such larger scales.

Researchers from the MIT Center for Collective Intelligence, for example, in [5, 21] discuss how technically very hard problems of global scale like climate warming can be solved by enabling collaboration among all those who are interested. Towards that, they created the Climate CoLab, an on-line platform for sharing ideas towards devising methods to enable climate changes. They report that CoLab has enabled large scale discussion among more than 10,000 members from more than 100 countries on over 400 proposals.

As a generic design consideration, it is recommended that DDSSP applications should have components which enable wider communication (and eventual collaboration) among all those who are associated with and interested in the problem.

4 Conclusion

Going by the trend evident from recent works [1, 11, 22, 23, 30–33], data-driven approaches using ITC technologies in conjunction with techniques from machine learning and statistics appear to play increasingly vital role in addressing societal

problems rooted in various socio-economic factors like poverty and lack of effective communication and coordination.

To address the grand challenges of complex social problems in low income geographies, this paper argues for increasing adoption of data-driven machine learning based solutions for enabling dynamic decision making towards determining timely preventive measures which can be applied by field agents of community outreach programs.

Paper identifies primary challenges in designing data-driven and machine learning based solutions, which can be deployed using ITC technology on larger scales. Some of the challenges are common to what any ML data-scientists would encounter, while others are relatively more specific to societal problems, like, difficulty in capturing temporally relevant semantic clues from the context and scaling ML models across sociocultural and geographical boundaries. To overcome these challenges, paper outlines series of design considerations, for example, design of continuous learning based predictive applications and structural analysis of data to enable fine grained analysis of local population a field agent is responsible for.

List of such design considerations arguably does not end here and should be augmented with additional design elements including enabling large scale data processing for wide scale adoption [11], collective collaboration, and techniques for knowledge graph generation and their use in deciding preventive measures.

References

1. Eggers, W.D.: *Delivering on Digital: The Innovators and Technologies that are Transforming Government*. RosettaBooks, New York (2016)
2. UNICEF India: What we Do? <http://unicef.in/Whatwedo>
3. CINI in India. <http://www.cini-india.org/>
4. Arjunpuri, C.: India Faces Epidemic of Missing Children. <http://www.aljazeera.com/indepth/features/2013/02/2013219121326666148.html>
5. Introne, J., Laubacher, R., Olson, G., Malone, T.: Solving wicked social problems with socio-computational systems. *KI-Künstliche Intelligenz* **27**(1), 45–52 (2013)
6. Schoder, D., Putzke, J., Metaxas, P.T., Gloor, P.A., Fischbach, K.: Information systems for “wicked problems” – proposing research at the intersection of social media and collective intelligence. *Bus. Inf. Syst. Eng.* **6**(1), 3–10 (2014)
7. Podder, S., Misra, J., Kumaresan, S., Dubash, N., Bhattacharya, I.: Designing intelligent automation based solutions for complex social problems. arXiv preprint [arXiv:1606.05275](https://arxiv.org/abs/1606.05275)
8. Miller, J.H., Page, S.E.: *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton University Press, Princeton (2009)
9. Ghosh, C., Paul, S., Kuntagod, N.S., Maitra, A.: Mitigating vulnerability of adolescent girls via innovative usage of digital technologies: insights from a field trial. In: *Proceedings of the 2nd International Conference on Science and Social Research, Malaysia* (2015)
10. Ethem, E.: *Introduction to Machine Learning*. MIT Press, Cambridge (2014)
11. Coulton, C.J., Goerge, R., Putnam-Hornstein, E., de Haan, B.: Harnessing big data for social good: a grand challenge for social work. *Am. Acad. Soc. Work Soc. Welf.* **11**, 1–21 (2015)
12. I. E. Agency: Electricity Access Database. <https://www.iea.org/energyaccess/database/>

13. Imbens, G.W., Rubin, D.B.: *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge University Press, Cambridge (2015)
14. Pearl, J., Glymour, M., Jewell, N.P.: *Causal Inference in Statistics: A Primer*. Wiley, Hoboken (2016)
15. Austin, P.C.: An introduction to propensity score methods for reducing the effects of confounding in observational studies. *Multivar. Behav. Res.* **46**(3), 399–424 (2011)
16. Haushofer, J., Shapiro, J.: The short-term impact of unconditional cash transfers to the poor: experimental evidence from Kenya. *Q. J. Econ.* **131**, 1973–2042 (2016)
17. Nichols, A.: Causal inference with observational data. *Stata J.* **7**(4), 507 (2007)
18. Mooij, J.M., Peters, J., Janzing, D., Zscheischler, J., Schölkopf, B.: Distinguishing cause from effect using observational data: methods and benchmarks. *J. Mach. Learn. Res.* **17**(32), 1–102 (2016)
19. Tikka, S.: Deriving Expressions of Joint Interventional Distributions and Transport Formulas in Causal Models. R-Package ‘causaleffect’ (2016). <https://rdrr.io/cran/causaleffect/>
20. DataKind: Learning From Text Messages to Help Teens in Crisis. <http://www.datakind.org/projects/learning-from-text-messages-to-help-teens-in-crisis>
21. Malone, T.W., Laubacher, R., Fisher, L.: How millions of people can help solve climate change. *Nova Next*. <http://www.pbs.org/wgbh/nova/next/earth/crowdsourcing-climate-change-solutions/>
22. Blumenstock, J.E.: Fighting poverty with data. *Science* **353**(6301), 753–754 (2016)
23. Barlow, M.: *Data and Social Good*. O’Reilly Media, Newton (2015)
24. Forman, G., Cohen, I.: Learning from little: comparison of classifiers given little training. In: Boulicaut, J.-F., Esposito, F., Giannotti, F., Pedreschi, D. (eds.) *PKDD 2004. LNCS (LNAI)*, vol. 3202, pp. 161–172. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-30116-5_17
25. Aste, M., et al.: Techniques for dealing with incomplete data: a tutorial and survey. *Pattern Anal. Appl.* **18**(1), 1–29 (2015)
26. Corani, G., Zaffalon, M.: Learning reliable classifiers from small or incomplete data sets: the naive credal classifier 2. *J. Mach. Learn. Res.* **9**(1), 581–621 (2008)
27. Silver, D.L., Yang, Q., Li, L.: Lifelong machine learning systems: beyond learning algorithms. In: *AAAI Spring Symposium: Lifelong Machine Learning*, vol. 13 (2013)
28. Davis, E., Marcus, G.: Commonsense reasoning and commonsense knowledge in artificial intelligence. *Commun. ACM* **58**(9), 92–103 (2015)
29. Lake, B.M., et al.: Building machines that learn and think like people. *Behav. Brain Sci.* **40**, e253 (2017)
30. Conference Series: International Conference on Information and Communication Technologies and Development. 2019: <https://www.ictdx.org/>, 2017: ictd2017.itu.edu.pk/, 2016: <https://ictd2016.info/>
31. AAAI-17 Workshop on AI and OR for Social Good (AIORSocGood-17). <https://sites.google.com/site/aiorsocgood17/>
32. The AAAI 2017 Spring Symposium on AI for Social Good (AISOC). <http://scf.usc.edu/~amulyaya/AISOC17/>
33. AI for Good Global Summit 2018. <https://www.itu.int/en/ITU-T/AI/2018/Pages/default.aspx>