Transdisciplinary Engineering for Complex Socio-technical Systems
K. Hiekata et al. (Eds.)

doi:10.3233/ATDE190137

Teaching Engineers Diplomacy, and Other Lessons for Machine Learning

Barrett S. CALDWELL[1]
Purdue University, West Lafayette, USA

Abstract. Advancements in capabilities for machine learning and artificial intelligence (MLAI) has led to growing questions and challenges for effective sociotechnical MLAI applications. It is important to consider applications outside of consumer product development, particularly in the realms of local, regional, national, and international policy development. There are intriguing opportunities to apply MLAI techniques to improving the efficiency of standard government ("bureaucratic") activities, including policy and diplomatic operations. Growing sociotechnical concerns address the limitations of applying machine learning tools to digital content, due in part to the unintentional (or even explicit) embedding of bias and prejudice into MLAI algorithms that become more difficult for others (particularly those without computer science training) to detect, correct, or reverse. The author's experience provides context to consider issues of MLAI involvement in policy creation in a more subtle operational context. Many MLAI applications are built on the analysis of an existing corpus of outcome products; this reasoning might be applied to the analysis of international policy and standards documents. This paper addresses two challenges to that approach. One is the difference between engineers and policymakers on the nature of debate, evidence, and conflict resolution. A second difference addresses, from the author's experience, emphasis on informal and interim processes, rather than final products, in the development of mutually agreed outcomes in national and international policymaking efforts.

Keywords. Policymaking, Foreign Affairs, Expertise, Cultural Context, STEM Research

Introduction

This paper grows out of the author's experience involved in a unique experience applying science, technology, engineering, and mathematics (STEM) research expertise to the realm of diplomacy and foreign policy. The United States Department of State, acting as the ministry-level government agency representing the United States' foreign affairs activities, is engaged globally to address sociotechnical issues affecting national and international interests. However, only a small fraction of the diplomatic officers have advanced degree training in STEM disciplines. Further, most tenured faculty in U.S. research-intensive universities focus on the research elements of their disciplines in conversation with other domain-based scholarly experts. A gap exists between the knowledge base of these STEM experts and the understanding and appreciation by the public of the importance of STEM research applied to societal concerns, including due to limited or poor communication by the STEM experts to the public [1].

[1] Corresponding Author, Email: bscaldwell@purdue.edu.

In an attempt to better utilize the STEM research expertise of U.S. universities, the U.S. Secretary of State Colin Powell announced the creation of the Jefferson Science Fellowship program (named for the first U.S. Secretary of State and third U.S. President, Thomas Jefferson) in 2003 [2-3]. The goal of the program from its origin was focused on "the need for creating new intersections between the culture of academic scientists and that of policymakers" [3, pg. 2]. Originally, program expenses were paid by major philanthropic agencies (the Carnegie Corporation of New York and the John D. and Catherine T. MacArthur Foundation); beginning in 2008, program expenses were assumed by the Department of State, with selection of Fellows managed by the U.S. National Academies of Science, Engineering, and Medicine. Most importantly, the salaries of Jefferson Science Fellows are paid by the faculty's home institutions, representing a critcial commitment to faculty contributions to national service by the participating universities. An additional, and not-to-be-underestimated, commitment is made by faculty, who willingly shift their career trajectories in order to participate the in the program. Of course, there is a significant shift in professional emphasis from the Fellow's domain of STEM expertise to the efforts of international policy, as well as the communication and translation of scientific and engineering research into the domains of diplomacy and geopolitical interactions. The shift is also physical and geographical, requiring the Fellow to operate "in residence" in Washington, DC for a calendar year (usually from August to August), with little or no time spent in traditional research, graduate advising, or scientific and professional conference participation.

From the first class of five Jefferson Science Fellows in 2004, the program has grown to include over 150 STEM faculty by 2019, in groups of 5-12 Fellows per year. The home disciplines of Fellows range across the physical and life sciences, agriculture, and engineering. Most Fellows apply their skills within their primary scientific domain in a policy bureau or other focused program emphasis in the U.S. Department of State or Agency for International Development. A few Fellows, by contrast, apply their expertise more broadly *across* STEM domains with a particular geographic focus, corresponding to regional bureau and country-specific offices within the Department of State. This paper is based on the author's experience as a Jefferson Science Fellow assigned to the "country desk" focused on the U.S. – Japan alliance within the East Asian and Pacific Regional Bureau of the U.S. Department of State, which included a variety of STEM discipline emphases in the areas of Environment, Science, Technology, and Health (ESTH). (NOTE: no comments in this paper are intended to represent official positions of the U.S. Department of State or any other government agency.)

1. An Engineer Goes Back to School

One of the most significant challenges for an established subject matter expert (SME) is to leave an environment where they are an acknowledge (and culturally validated) authority and adapt to a new setting outside of their area of expertise. In August 2016, soon after moving into a new apartment, the author began a one-week "new employee orientation" including such presentations as "Washington Statecraft," designed for those who had not previously worked as managers within the main Department of State environment. In essence, the experience is similar to that of a new graduate (or even one-year exchange) student attempting to begin a new program, learning what terms and practices from their prior education will be applicable to a new environment. (As a senior faculty member, though, it had been over 30 years since the author's prior experience as a student on a new campus.)

This shift in SME orientation represents important distincitions in process and content orientations of knowledge management, described as internalisation, socialisation, externalisation, and combination [4]. In addition, the SME has competing challenges to recognize his/her own individually configured and tacit (socially configured as a member of a professional community of practice) knowledge bases to make them explicit, and to reinterpret new explicit information in previously unfamiliar tacit contexts [1,4]. In some cases, the important contribution of the SME is not to provide detailed knowledge about the specifics of the domain, but to identify the context of application, or communicate effectively regarding basic elements of the subject to intelligent non-experts. These skills have been desribed as distinct dimensions of expertise, that are not simply aggregations of SME within a STEM domain [5].

The volume and quantity of communication was, for this author, a very distinctive change. A conference or journal abstract in a professional publication is often between 100 and 250 words (the length of an abstract for this Transdisciplinary Engineering conference); these are considered "brief summaries" of the SME contributions of the paper. By contrast, a foreign affairs officer for a regional bureau may be required to select no more than three items for summary for the daily updates of an entire nation's ESTH or political activity, where each item is no more than five lines (roughly 50 words) and does not contain detailed technical reference terminology. These approaches to summarizing information require a different integration of explicit and tacit information and context. Non-academic (corporate, diplomatic, economic or military) audiences may be more focused on a "BLUF" (bottom line, up-front) model of communication [6], combined with this extreme discipline of length, to support quick knowledge management and updating for senior principals managing overall broad geographical regions and diplomatic concerns.

Coincidentally, one of the ESTH issues assigned to the author during his Jefferson Science Fellowship tenure included the broad domain of machine learning and artificial intelligence (MLAI). From an economic and employment perspective, there are many elements of MLAI of direct interest to U.S. and Japan from industrial, societal, and sociotechnological perspectives; advances in artificial intelligence for precision medicine, disaster response, and support for an aging society have all been considered important facets of Japan's Society 5.0 initiatives. There was, however, an additional query from the governmental employees involved in MLAI initiatives: how can we use such tools to improve the administrative (bureaucratic) practices of foreign affairs and international policy? The remainder of this paper considers four possible lessons from the author's experience, framed as considerations that influence or limit the application of more general MLAI approaches in the service of diplomacy.

2. Learning as Context and Relationship Transfer

Within six weeks of the initial assignment on the "Japan Desk," the author was involved in at least three distinct ESTH projects in different areas of STEM research specialization: precision medicine (for the Cancer Moonshot program), artificial intelligence and societal applications (for the Science and Technology in Society

Forum), spaceflight applications (for JAXA satellite launches). These projects were all in addition to overall monitoring of economic and political elements of U.S. – Japan relations. In this environment, the most familiar feature elements (e.g., details of metadata big data analytics; orbital coverage of a positioning satellite) for a STEM researcher were not the critical components of developing skill in the diplomatic "arts".

In machine learning, detection and extraction of features in the physical environment based on geometric or other mathematical analysis is a very common operational approach. Overall pattern and context transfer from initial training sets to novel situations, by contrast, remains a daunting effort. Understanding broader societal elements of Japanese culture were fostered and explicitly encouraged through considering two qualitative works: *The Chrysanthemum and The Sword* (historical anthropology) [7] *Shogun* (historical fiction) [8]. The value in these works, at least to the author, was in a broader gestalt contemplation of context, especially in terms of the considerations of a fundamentally distinct culture. Since the Jefferson Science Fellowship is in itself an exploration of a fundamentally distinct culture between STEM research and policymaking, the appreciation of context shift was more evident than for someone who might consider implementations of machine learning in the very different domain of diplomacy as simply an instantiation of a very familiar tool to a novel application.

3. Training Sets are Not Objective, Unbiased Data

From a technical perspective, the need for a substantial training set is a *sine qua non* of any successful machine learning algorithm: the software must be provided with unambiguous examples of a object (an image of a face or representation of an animal; a letter in several fonts or an accurate word translation; a phrase of music or the specific overtones of an individual instrument or singer) in order to generalize examples of the object across new presentations (the test set). Of course, for these machine learning approaches to work, there must be a suitable corpus of previously identified examples, which qualify as the training set.

The question of how many examples is required for a suitably confident machine learning prediction applied to new observations is a challenging statistical one, both in terms of the number of training set examples and the size of the training set compared to the test set. No single answer exists for either total examples or extrapolation (although many researchers may use training sets several times larger than the test set); the expense of human validation and annotation of training set data becomes an operational (rather than theoretical) limitation [9].

Further, there are increasing acknowledgements that online, web-based training set data cannot itself be assumed to be an objective, unbiased source of information when developing and implementing policies on a national and global scale. At the most basic level, the history of available online information is overwhelmingly English. Although the percentages of web content (not to mention social media content) in English have been dropping over the 21st Century from over 80% to 30-50%, the percentages of internet content still drastically oversample the percentages of the world's population who speak English well (approximately 20%) or natively (approximately 5%) [10-12]. A recent problem of bias in image recognition was shown by the American Civil Liberties Union, who noted that a number of members of the U.S. Congress, particularly women and persons of color, were misidentified as matching photos from arrest records rather than their status as political representatives (or even political leaders) [13]. As a result, machine learning training sets from general internet archives cannot be assumed to manage unbiased classification of even “objective” items such as official photo images from government sources; these concerns become even more critical as other cognitive bias elements of cultural discrimination or marginalization are seen to have impacts on policy development and decision making [13-14].

4. Team Drafts, not Final Authored Documents, are the Relevant Corpus

As described above, the use of a poorly calibrated or unvalidated training set can result in even more devestating consequences for machine learning applications to policymaking than most non computer scientists may understand. The choice of a suitable training set, and appropriate curation of those data, becomes an important policy decision in itself. Policy documents have an additional burden for curation, as the culture of “statecraft” this author experienced was quite distinct from the academic culture of the university.

One of the most basic differences observed between academic research writing, and the development of policy statements presented by senior diplomats, is the attribution of authorship. A search through ResearchGate or Google Scholar can identify articles written by a particular researcher, and determine areas of expertise and relative contributions based on the documents that have been published under that author’s name. Professional ethics also determine how or whether order of authorship maps onto relative contribution of the authors. Order of authorship might even change based on new contributions from one author from initial draft to final submission; only the content and the authorship of the final submission is noted in the publication record.

However, the corpus of training set documents for policy making does not represent a similar mapping of attribution. Even if a senior diplomat or political leader (“principal”) delivers a speech on a particular topic, it would be considered very unlikely that the principal who delivered the speech studied all of the background material and conducted the relevant scholarly inquiry themselves. In other words, the major effort of writing a policy document may have been contributed by members of a team – or, in the sake of a policy with multiple disciplinary or regional “equities” of interest, multiple teams. The non-academic setting of policymaking may completely obscure or diminish the concept of individual attribution, thus limiting the mapping of relative contributions to specific subject matter experts. In addition, the tracking of only final documents will certainly mask the negotiation and nuance associated with the intentions or wordings suggested in prior drafts.

A distinct metaphoric approach to considering machine learning of evolving policy documents might include “change detection” approaches, such as those used in video image compression. If video is considered as a series of sequential images, change detection represents the analysis of how visual elements change from one image to the next one in sequence, without attempts to determine or identify the specific objects represented in the images [15]. A second possible conceptual alternative might be one of social network and influence analysis, an element of social science study of information flow that focuses more on interaction processes rather than specific outcomes. Email metadata provide the “To:” and “CC:” lines of messages that can be analyzed, even for email that might not be considered public for reporting on the draft discussions. Considering who is “looped in” on drafts for consideration and consultation, in addition to the evolution of text from initial concept to formal declaration, helps to document how equities and stakeholders were considered in policy formulation.

5. Data, Debate, and Advice

Another shift in context emphasis from the academic STEM environment and the diplomatic corps involves the type and use of expertise involved in policymaking. Even the "currency" of data-driven, vs. persuasive, knowledge sharing and influence may have significant gaps between the academic scientist and the policymaking principal.

A STEM researcher participating as a Jefferson Science Fellow may have spent 25 years or more perfecting the technical and presentation skills to create a journal manuscript or conference proceedings document according to the recognized standards of the local research community and professional discipline. Quite simply, it becomes difficult to remember what it is like *not to know how* to write such a technical paper. In addition, since the supposed purpose of such papers is the clear and effective presentation of data and logic arguments, it is easy to assume that the researcher's job of communication is completed by simply presenting the data, research analysis, and conclusions, using a "truth wins" philosophy. For the researcher, the clear presentation of assumptions, methods, and empirical findings is how science progresses. By contrast, the researcher may be interpreted by others as saying, "of course it's obvious that this is the answer, and you must be ignorant to not agree with this clear presentation of empirical findings" (which are not clear to non-experts). These gaps can be suggested as primary failings on the part of research specialists when trying to communicate their findings to a broader (non-scientific) audience [1].

Those who play a more explicit role as bridging the STEM and policy worlds may focus efforts on determining how empirical data may enhance or restrict one's strategic options in a sociotechnical, and not merely analytical, context. A distinct approach to science and technology suggest that the value of advice comes from the ability to foresee and respond to a range of possible events and global sociotechnical challenges from a policy standpoint [16]. The challenge is not to force the principal to make a specific prediction about future events, but to consider and address issues of "deep uncertainty," where the overall range of policy options and decision responses may be difficult to anticipate [17]. The value of MLAI models and simulations in this "deep uncertainty" environment is not to generate specific optimal solutions, but to help provide recommendations to policymakers that expand their strategic and operational flexibility of response [16-17].

6. Summary of Impacts for Machine Learning Scholars

It is beyond the scope of this paper (and perhaps contrary to the goals of a Transdisciplinary Engineering conference) to focus or limit the discussion of advances in artificial intelligence applied to government activities to specific considerations of predictive accuracy or comprehensiveness of particular machine learning algorithms. It is important to recognize that over 60 years of research have continued to emphasize that effective translation between natural human languages is a function of recognizing context, intent, and nuance rather than pattern matching of archetypal objects or explicit denotations of individual words. Machine learning algorithms remain challenged by such complexities in meaning and strategy; in fact, most humans have distinct limitations in these skills without considerable language fluency practice and experience. The author's experience with international science and technology policy suggests not a specific set of deep learning mechanisms to optimize, but a recognition of the barriers of entire classes of machine learning approaches.

Although there is a common set of assumptions among scholars and inventors in the machine learning and data analytics domains regarding the objective and unquestioned status of training set data, a number of recent experiences in machine learning indicate that scholars must take considerable care in recognizing the limitations and subjective natures of text documents making up the diplomatic corpus [10-14]. Since the process of diplomacy and international policy negotiations is rarely complete (or rarely based on independent observations), change detection regarding evolving drafts (including the reasons why particular language was considered and discarded) is a more appropriate model for applying machine learning algorithms than explicit matching of assumed objects or patterns [15].

7. Conclusion

After 25 years as an engineering research faculty member, the author experienced a unique opportunity to shift cognitive and sociotechnical contexts to go “back to school” and experience the world of international science and technology policymaking. In addition to important personal lessons in understanding major shifts in context and expertise comprehension, the author's experience provides metaphors to consider opportunities to apply machine learning and artificial intelligence (MLAI) techniques to the world of statecraft. Instead of thinking of policy documents and agreements as simply another application of known MLAI techniques, the author's experience helped to underscore the importance of context transfer and subtly negotiated understanding of the “other” in human-human interactions. This paper does not provide specific recommendations for advances to algorithms to improve bureaucratic or strategic operations in foreign affairs. However, this paper does attempt to address critical sociotechnical elements of bias and conflicting assumptions on the nature of a good operational solution to an ill-defined systems engineering problem.

The opportunity to engage in foreign affairs applications of environment, science, technology and health (ESTH) domain experience is a unique form of public communication of science and engineering in a societal context. STEM researchers who can work to address ESTH issues in a cooperative and supportive advising role, rather than as an explicitly attributed expert, will experience distinct cultural shifts in their domains and applications of expertise. However, the potential impact of their knowledge can reach beyond their specific areas of research expertise and improve the availability of knowledge to respond to critical and dynamic policy challenges.

Acknowledgement

The author expresses his appreciation to his colleagues at the U.S. Department of State, and the Jefferson Science Fellowship program (managed by the U.S. National

Academies of Science, Engineering, and Medicine). The opinions expressed here are the author's, and do not reflect the official perspectives or positions of any government agency.

References

[1] B.S. Caldwell, Spaceflight-relevant stem education and outreach: Social goals and priorities, *Acta Astronautica,* vol. 112, 2015, pp. 174-181.

[2] MacArthur Foundation, *MacArthur Supports New Science and Security Fellowship Program at U.S. Department of State*, press release 2003, Accessed 14.03.2019. [Online] Available: https://www.macfound.org/press/press-releases/macarthur-supports-new-science-and-security-fellowship-program-at-us-department-of-state-october-8-2002/

[3] Carnegie Corporation of New York, Carnegie Results Summer 2010: Jefferson Science Fellows. Accessed 14.03.2019. [Online] Available: https://www.issuelab.org/resources/8584/8584.pdf

[4] J. Guinery, Capturing decision input content to support work system design, *Proceedings of the Tenth International Symposium on Human Factors in Organizational Design and Management*, Grahamstown, South Africa, 2011, pp. 259-265.

[5] S.K. Garrett, B.S. Caldwell, E.C. Harris and M.C. Gonzalez, Six dimensions of expertise: a more comprehensive definition of cognitive expertise for team coordination, *Theoretical Issues in Ergonomics Science*, vol 10, 2008, pp. 93-105.

[6] K. Sehgal, 2016, *How to write with military precision*, Accessed 14.03.2019. [Online]. Available: https://hbr.org/2016/11/how-to-write-email-with-military-precision

[7] R. Benedict, *The chrysanthemum and the sword: Patterns of Japanese culture*. Houghton Mifflin Harcourt, Boston, 1946.

[8] J. Clavell, *Shogun.* Dell, New York, 1975.

[9] R. L. Figueroa, Q. Zeng-Treitler, S. Kandula, and L. H. Nog, Predicting sample size required for classification performance, *BMC Medical Informatics and Decision Making*, vol 12, 2012. [Online]. Available: https://bmcmedinformdecismak.biomedcentral.com/articles/10.1186/1472-6947-12-8

[10] Unbabel, 2015, *Top Languages of the Internet, Today and Tomorrow*, Accessed: 18.03.2019. [Online]. Available: https://unbabel.com/blog/top-languages-of-the-internet/

[11] H. Young, n.d., *The digital language divide: How does the language you speak influence your experience of the internet?* Accessed 18.03.2019. [Online.] Available: http://labs.theguardian.com/digital-language-divide/

[12] Internet World Stats, 2018, *Internet World Users By Language,* Accessed 18.03.2019. [Online.] Available: https://www.internetworldstats.com/stats7.htm

[13] M. Paolini-Subramanya, 2018, *Facial Recognition, and Bias*, Accessed: 18.03.2019. [Online.] Available: https://medium.com/datadriveninvestor/facial-recognition-and-bias-5bd963d83813

[14] M. McElaney, 2018, *Cognitive Bias in Machine Learning*, Accessed: 18.03.2019. [Online.] Available: https://medium.com/ibm-watson-data-lab/cognitive-bias-in-machine-learning-d287838eeb4b

[15] B. Shahraray, Scene changes detection and content-based sampling of video sequences, In: Proceedings SPIE 2419, Digital Video Compression: Algorithms and Technologies, San Jose, CA, 1995, https://doi.org: 10.1117/12.206348

[16] V. Turekian and T. Kishi, Science and technology advising in today's foreign policy, *Science and Diplomacy*, vol 6., 2017, Available: http://www.sciencediplomacy.org/perspective/2017/science-and-technology-advising-in-todays-foreign-policy

[17] B. S. Caldwell, Cognitive challenges to resilience dynamics in managing large-scale event response, *Journal of Cognitive Engineering and Decision Making,* vol 8, 2014, Available: https://journals.sagepub.com/doi/abs/10.1177/1555343414546220