Can Bots be Better Learners than Humans?

Wassim Derguech and Mathieu d'Aquin

Insight Centre for Data Analytics - National University of Ireland, Galway
firstname.lastname@insight-centre.org

Abstract. In this paper, we discuss how Learning Analytics, as the activity to capture and analyze people's learning behaviors in order to improve their learning experiences, could be used as a way for bots to "learn how to learn" and how this might have a greater impact than the apparent improvement it would enable for Artificial Intelligence. Through exploring this particular scenario as part and in the spirit of the *Recoding Black Mirror* workshop, we extrapolate a potential negative use of technologies currently being developed both in Technology Enhanced-Learning and in Artificial Intelligence to anticipate on some potential ethical issues they might generate, as a first step towards potentially more ethics-aware design and development activities in those areas.

Keywords: Learning analytics, Machine learning, Learning bots, Fake behavior

1 Introduction

In the last decade, many companies started to introduce support services using bots: "enhanced conversational agents that can chat with users"[1]. Indeed, Gartner predicts that, by 2019, 20% of user interactions with smartphones will take place using virtual personal assistants.[1] Bots use knowledge bases with either static or dynamic set of patterns to answer a query or maintain a conversation with a human user [2]. Knowledge bases can be enriched from existing conversations or new textual or multimedia resources to help the bot "learn" how to answer questions, propose a decision or compete with humans (e.g. bots in games). To a larger extent – as in episode 1 of the second season of the British science fiction anthology series *"Black Mirror"* (*"Be Right Back"*) where an artificial program is used to simulate the behavior of a deceased person from their online social interaction – one can expect bots to increasingly use the web as a learning platform to create the knowledge base for their artificial intelligence.[2]

In addition, there is currently a trend in using technological means to improve the ability of human users of the Web to learn using online platforms and tools.[3] *Learning Analytics* [3] encapsulates the idea of monitoring, analyzing and assessing the behavior of learners to support them being more efficient.

[1] http://www.gartner.com/newsroom/id/3551217
[2] With some unsuccessful attempts: e.g., https://techcrunch.com/2016/03/24/microsoft-silences-its-new-a-i-bot-tay-after-twitter-users-teach-it-racism/
[3] See as a state of the art example the AFEL project: http://afel-project.eu

In this paper, we therefore ask the questions: "Can bots learn from humans how to be efficient at learning?". More precisely, we discuss the idea that online bots could use the Learning Analytics from humans to target their own learning towards a particular topic and become themselves more effective in integrating such a domain. While this could appear altogether as a positive side effect of Learning Analytics for Machine Learning and Artificial Intelligence, we also discuss the possible pitfalls of enabling bots to simulate the learning behavior of humans from traces and data that are, after all, only capturing what is happening at the surface of the cognitive process of learning.

2 Learning Analytics: Basic Concepts

Learning analytics is defined in [3] as "the measurement, collection, analysis and reporting of data about learners and their contexts, for purposes of understanding and optimising learning and the environments in which it occurs". It has grown very rapidly in recent years as an academic research area at the intersection of data analytics and education science, but also in practice with many universities adopting learning dashboards and recommender systems for purposes such as improving learning design (e.g., [4]) or retention (e.g., [5]).

In the simplest of cases, Learning Analytics can be seen as business intelligence for educational institutions. In such cases, institutions would generally collect information about each student, including demographics, background, entry-level information, and about the resources that they have at their disposal, including the eLearning system in place, the library, online courses, forums, etc. Crucially, they will also collect information about the activities of each student using those resources. The idea is, in many cases, to extract from those traces indicators that enable the institution to predict student's performance. The same analysis will often be used to identify effective learning practices and provide feedback to learning design activities. In the simplest form, this involves identifying modules or lecturers who appear to lead to greater students' success.

While what is described above represents one of the most common scenarios of Learning Analytics, a trend is currently emerging in enabling the use of Learning Analytics techniques not only for educational institutions, but also for the learners themselves [6]. Such scenarios are naturally more targeting self-directed learners who might use many different platforms and tools for learning online (MOOCs, open educational resources, social media). The activity data being captured and the way to analyze them therefore need to encompass a large part of all online activities of the learner [7]. Such approaches become naturally more oriented towards the social aspects of learning, through including in the learner's profile their social connections and context. It implies that a key part of the reflection towards one's own learning involves sharing and comparing one's learning activities and analytics with others.

3 Using Learning Analytics to "Learn How to Learn"

Many of the data analytics techniques used in Learning Analytics as described above are very much based on Machine Learning. The idea is to inspect large amounts of data about many "learning trajectories", and try to derive from that the indicators enabling to predict whether a particular trajectory is going to lead to success, using models which are more or less opaque.

In a very straightforward manner, we can envisage for bots to use the exact same models learned from humans' learning experiences to figure out what to do to achieve a similar goal. In other words, if we consider the objective of a bot to integrate knowledge from a certain domain in order to better deal with inquiries on that domain, or for any other task a bot might need to achieve, we can imagine the bot using the sum of human learning experiences that are captured in a Learning Analytics model to drive the choices it makes in gathering the knowledge it requires: Which resource to integrate, and which to ignore; how to assess the coverage of learning, its value and its bias; and how to use this information to further learn. Of course, this is assuming that technologies for knowledge extraction [8] are developed to the extent that, knowing which online resources would contain the knowledge required by the bot, it is possible to distill such a resource to obtain that knowledge in a form which the bot can exploit. Indeed, if we did assume that, the main challenge remaining would be in the discovery and selection of resources, which is what we suggest the use of Learning Analytics from human learning experiences would somehow solve.

Now, while a bot is predictable, as it is a simple application that acts based on the way it was programmed, humans are not. Humans can react differently to a similar situation while bots usually take the same action given the same inputs provided. Humans usually take a longer time to ingest data and appreciate the information they are receiving before taking any action, while bots react systematically, and therefore quickly [9].

This fundamental difference between bots and humans makes the learning trajectories of bots more "programmable". They can search data, analyze it, and use the power of computation to identify millions of fact-based options and rank the best one faster than a blink of a human eye. However, human emotions cannot be easily understood by bots through simply analysing the text or activities of human users. For a human learner, emotions are critical, including in problem solving, and performing tasks across domains, and has a considerable impact on learning performance [10]. Therefore, designing a bot that identifies and uses human emotions during their learning trajectories remains a research challenge.

4 A step too far?

One of the most obvious objections to the use of learning analytics and the development of bots as described above is that, by nature, Learning Analytics can only capture the cognitive process of learning from a surface point of view. Indeed, there have been many theories and models of learning proposed in cognitive psychology. Here, we refer to one of the most recent and more suitable to social environments associated with online learning: The co-evolution model [11]. In a nutshell, the idea of the co-evolution model is that learning happens as a side effect of the interaction between the cognitive system of the learner and the social system in which they operate. It is the result of the internalization of knowledge or aspects of knowledge that are encountered by the learner in their social environment, and that are directly or indirectly creating a friction with existing knowledge or aspects of knowledge that the learner has already internalized. It is new information or information presented through a point of view or a level of complexity that disrupt the cognitive system in a significant way. In turn, it is the process of externalizing one's views to the social system which is at the basis of knowledge creation.

Considering this model, a key to learning analytics as considered in the AFEL project is that it should essentially be about trying to assess what are the artifacts and conditions that best generate such a friction [7]. Indeed, through finding out which resources would introduce an increment in level of granularity or complexity with respect to what the learner might have seen, or which would introduce a new point of view on a potentially controversial topic, the idea is that the learner can be more focused in their learning.

Therefore, in the use of Human Learning Analytics as described in the previous section, the main objection considered can be better formulated as that, by only using the data that are surfaced through the technological mediation of learning (the platforms and tools), bots will by nature reduce the learning experience to a set of simplified, easily programmable aspects. They will reduce a complex set of decisions which might be driven by the complex relationships between the human learner, the knowledge and the resources in which it is encapsulated, which necessarily involve some degree of emotion, to over simplified indicators.

The most obvious side effect of this is a potential loss in the accuracy of the decisions taken (which, to a large extent, is a criticism that apply to Learning Analytics generally: that it can only be as accurate as what it can capture, which necessarily misses out on many aspects). However, there is a potentially more profound issue in this if we assume, as the current trend in applying Learning Analytics might suggest, that more and more educational institutions will rely on such approaches to data capture and modeling to understand, monitor and assess the learning of their students. Indeed, by reducing learning to something that can be understood from a purely analytical stand, we open the door to learning practices progressively being more and more targeted towards maximizing those analytical indicators, rather than actually promoting learning. This would primarily translate, first, into students aiming to "act" like good learners, realizing activities that they know will get them to be perceived as high performers (which is a known side-effect of introducing performance indicators in any domain [12]). Closer to the topic of this paper however, and in the spirit of extrapolating a potentially much more dystopian scenario, bots being as described in the previous section are better at being systematic in their behaviors and programmed to maximize such indicators, they might end-up being used as a very degraded form of “intelligent learning assistants”, helping learners in finding what to do to make themselves look good in the eyes of analytics. They might even be used to fake such learning behaviors on behalf of the student, making them become a part of the scenarios of ”Weapons of Math Destruction” [13] where the author highlights that technology might be used for encoding biases into algorithms.

Some elements of solutions might be found of course, including more robust identification mechanisms to ensure that what is analyzed is indeed the activity of the human learner, and not a bot working on their behalf, or as mentioned in the previous section, the integration of more sophisticated aspects related to human learning in the analytics process, including emotion analysis. However, without countering the current trend in attempting to understand human behaviors, in learning or in other domains, purely from the point of view of analytics, we take the risk to create measures that make bots take over human activities precisely because they are much less sophisticated about them. That would translate into a reduction in both our ability to perform such activities to the standards expected by analytics and benchmarked by the bots, and in the actual quality of the activity, including learning, which would become reduced to its most shallow expression.[4]

5 Conclusion

In this paper, we presented the general idea that, as bots are learning from human behaviors and learn to emulate them (as in *Black Mirror* Episode “Be Right Back”), they might as well learn from humans’ learning behaviors through Learning Analytics. However, we also discuss how, in learning or in other areas, relying on the shallow traces of human behaviors, bots can only learn at the surface. They can learn to fake the human behavior and therefore, might be trivially turned into assistants to help humans in automatizing the faking of a positively viewed behavior such as learning. This is of course very related to a common criticism of Artificial Intelligence, which is that it cannot actually create intelligence, only simulate the appearance of intelligence (see for example, the “Chinese Room Argument” [14]). However, what is discussed here goes a step further: By making it a trend to assess human behaviors from necessarily shallow digital traces, we risk to restrict the goal of Artificial Intelligence to only being a simulation of the appearance of intelligence, rather than seeing it as a limitation of the technology. There are already much discussions on how this confusion between appearances and actual behaviors often materialize in sometimes high-level societal situations with humans (entertainment, politics, etc.), so it would truly be a shame if we ended up, on purpose, making our bots as shallow as we sometimes are.

[4] An interesting parallel here can be drawn with search engine ranking algorithms leading to bots being used for Search Engine Optimization (SEO), and not really to better online content.

Acknowledgement

This work has received funding from the European Union's Horizon 2020 research and innovation programme as part of the AFEL (Analytics for Everyday Learning) project under grant agreement No 687916.

References

1. Klopfenstein, L.C., Delpriori, S., Malatini, S., Bogliolo, A.: The rise of bots: A survey of conversational interfaces, patterns, and paradigms. In: Proceedings of the 2017 Conference on Designing Interactive Systems. DIS '17, New York, NY, USA, ACM (2017) 555–565
2. Reshmi, S., Balakrishnan, K.: Implementation of an inquisitive chatbot for database supported knowledge bases. Sādhanā **41**(10) (Oct 2016) 1173–1178
3. Ferguson, R.: Learning analytics: drivers, developments and challenges. International Journal of Technology Enhanced Learning **4**(5-6) (2012) 304–317
4. Lockyer, L., Heathcote, E., Dawson, S.: Informing pedagogical action: Aligning learning analytics with learning design. American Behavioral Scientist **57**(10) (2013) 1439–1459
5. Dietz-Uhler, B., Hurn, J.E.: Using learning analytics to predict (and improve) student success: A faculty perspective. Journal of Interactive Online Learning **12**(1) (2013) 17–26
6. Bull, S., Ginon, B., Kay, J., Kickmeier-Rust, M., Johnson, M.D.: LAL workshop: learning analytics for learners. In: Proceedings of the Sixth International Conference on Learning Analytics & Knowledge, ACM (2016) 496–497
7. d'Aquin, M., Adamou, A., Dietze, S., Fetahu, B., Gadiraju, U., Hasani-Mavriqi, I., Holtz, P., Kimmerle, J., Kowald, D., Lex, E., Sola, S.L., Maturana, R.A., Sabol, V., Troullinou, P., Veas, E.: AFEL: Towards measuring online activities contributions to self-directed learning. In: Proceedings of the EC-TEL 2017 workshop ARTEL: Awareness and reflection technology enhanced learning. (2017)
8. Gangemi, A.: A comparison of knowledge extraction tools for the semantic web. In: Extended Semantic Web Conference, Springer (2013) 351–366
9. Gilani, Z., Farahbakhsh, R., Tyson, G., Wang, L., Crowcroft, J.: An in-depth characterisation of bots and humans on twitter. CoRR **abs/1704.01508** (2017)
10. Azevedo, R., Millar, G.C., Taub, M., Mudrick, N.V., Bradbury, A.E., Price, M.J.: Using data visualizations to foster emotion regulation during self-regulated learning with advanced learning technologies: A conceptual framework. In: Proceedings of the Seventh International Learning Analytics & Knowledge Conference. LAK '17, New York, NY, USA, ACM (2017) 444–448
11. Kimmerle, J., Moskaliuk, J., Oeberst, A., Cress, U.: Learning and collective knowledge construction with social media: A process-oriented perspective. Educational Psychologist (50) (2015)
12. Fortuin, L.: Performance indicatorswhy, where and how? European journal of operational research **34**(1) (1988) 1–9
13. O'Neil, C.: Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. Crown Publishing Group, New York, NY, USA (2016)
14. Searle, J.: Chinese room argument, the. Encyclopedia of cognitive science (2001)