

A Neural-Symbolic Approach for User Mental Modeling: A Step Towards Building Exchangeable Identities

Hossein Rahnama^{a,c,d}, Marjan Alirezaie^{b,c} and Alex Pentland^d

^aRyerson University, RTA School of Media, Toronto, Canada

^bCentre for Applied Autonomous Sensor Systems (AASS), Örebro University, Sweden

^cFlybits Inc.

^dMIT Media Lab, Massachusetts Institute of Technology

Abstract

Combining symbolic-reasoning and data learning in a unified double-loop learning system can contribute to the emergence of artificial intelligence solutions that are more adaptive to social and behavioural context. This paper presents a hybrid user modeling framework that relies on the integration of machine learning and reasoning methods equipped with formally represented domain knowledge. We find that this approach contributes to the design of context-aware systems that require less data, manage bias better, provide better transparency and can handle data sparsity more effectively. We present the impact of our work in different social domains from building trusted digital surrogates to decentralization of social recommendation services. Our approach can construct software agents from identity and expertise of users and allows such entities to become more digitally portable. Our approach also contributes to the emergence of expertise sharing paradigms that are less prone to biases and more privacy preserving. The paper uses these domain applications to validate the scalability and versatility of our approach augmented with principles of open and transparent algorithms.

Keywords

Neural-Symbolic Integration, Exchangeable Identities, User Mental Modeling

1. Introduction

In many of recent applications of human-computer interaction (HCI), a machine enabled with artificial intelligence (AI), receives data from its human user or the environment in order to first learn the user behaviour and then adapt its services to become more compatible with the user's needs or preferences (a.k.a, personalization). In other words, in classic HCI models (see Figure 1) the source" user, whose data is employed to train the machine, and the end" user, who is served by the machine, are identical [1].

However, separating the source and the target users (see Figure 2) can extend the spectrum of HCI applications towards scenarios relying on what we call *borrowable identity*. For instance, imagine a corporate lawyer who would like to provide her expertise to a network of clients through an intelligent system which is equipped with a computational model reflecting the lawyer occupational knowledge as well as her identity. Given a comprehensive and adaptive user model, the intelligent machine will ideally respond to the clients with a high degree of relevance and also a reduced cost compared to the lawyer in-person rate sheet¹. In this case, the clients (i.e., target users) will have the ability

In A. Martin, K. Hinkelmann, H.-G. Fill, A. Gerber, D. Lenat, R. Stolle, F. van Harmelen (Eds.), Proceedings of the AAAI 2021 Spring Symposium on Combining Machine Learning and Knowledge Engineering (AAAI-MAKE 2021) - Stanford University, Palo Alto, California, USA, March 22-24, 2021.

EMAIL: rahnama@mit.edu (H. Rahnama); marjan.alirezaie@oru.se (M. Alirezaie); sandy@media.mit.edu (A. Pentland)

ORCID: 0000-0002-4001-2087 (M. Alirezaie); 0000-0002-8053-9983 (A. Pentland)



© 2021 Copyright for this paper by its authors.

Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

¹<https://www.media.mit.edu/projects/augmented-eternity/overview/>




Figure 1: HCI user model with similar source and target user.




Figure 2: HCI user model with distinct source and target users.

to *borrow the identity* of the lawyer (i.e., source user). The applications of the extended HCI model presumed to hold the mental model of the user, can potentially go beyond the borrowable or digital identity to what we call *swappable identities* that let one understand how reality appears to others with different mind-set, or simply to see the world through different lenses. In Section 4, we further discuss possible applications of the extended user model that may be useful for individuals, communities and corporations.

In order to achieve a borrowable (or digital) identity as a computational model of the (source) user, a machine needs to go beyond gaining knowledge about what the user likes or dislikes (e.g., in recommendation systems), and attain the knowledge of her mentality, i.e., way of thinking, as well as her personality, i.e., way of acting [2]. There are rich literature resources in cognitive science and psychology with the focus on development of mental models in humans. By mental model we refer to a knowledge structure that represents someone’s intuitive perception of her environment, the relationships between different entities in the environment and also her way of thinking or reasoning upon the perceived world [3]. More specifically, a mental model in an agent (either biological or artificial) shapes its behaviour and plays a paramount role in the reasoning and decision-making of the agent [2].

The proposed idea in this position paper is focused on enabling machines to build the digital mental model of humans based on aggregated digital footprints. A computational agent equipped with such a model can potentially become a reliable digital substitution of its user in serving and interacting with others in different possible situations. Mental models in humans are constructed incrementally as the result of a learning process called double-loop learning [4] shown in Figure 3. According to double-loop learning (vs. single-loop learning), construction of a mental model is performed within an endless loop between two processes of updating and extending the contents of the model based on the feedback from the environment and making decisions based on the updated model [4, 5].




Figure 3: Double-loop learning model: While the mental model is being updated during the learning process it contributes to the decision making process.

Recent successes in advanced machine learning (ML) methods have led to a proliferation of user modeling applications (e.g., recommendation systems or companion technologies) [6, 7]. However,

building a user model solely based on data-driven methods is associated with a number of challenges including (1) dependency on large amount of data(2) data bias(3) lack of transparencyand (4) data sparsity [8]. Due to the aforementioned issues, it is often the task of human to interpret the results of such methods for purposes such as transparency assesment and knowledge generation [9]. On the other hand, in double-loop learning, constrution (or updating) of mental models also relies on a knowledge generation process that translates the feedback from the environment (learning outputs) to meaningful knowledge.

According to recent research works, enabling ML methods with reasoning abilities can significantly improve the results by addressing the aforementioned issues [10]. More specifically, the idea behind combining learning and reasoning methods is to get the benefits of both sides: robust learning methods (mainly subsymbolic ones) as well as interpretability of symbolic representation and reasoning approaches [11]. Furthermore, such hybrid methods can also contribute to automate both knowledge generation and decision making processes within a double-loop learning model.

In this paper, we propose a user mental modeling approach that, as depicted in Figure 4, relies on an integration of machine learning and reasoning methods, equipped with formally represented domain knowledge. As we can see, the proposed model complies with double-loop learning in the sense that: (i) learning is always associated with the current contents of the mental model, and (ii) the mental model, per se, incrementally grows and is updated based on the output of the integrated learning-reasoning process applied upon newly received data. As we will discuss in the following sections, the proposed approach contributes to build a multi-modal mental model of a human user based on personal characteristics latent in her digital footprints. Due to its basis relying on theories in psychology, the proposed model has the potential to become the digital identity of the user and represent how the user will observe, think and (re)act to changes in a given environment.

It is our goal to specifically target symbolic representation and reasoning methods which (a) cover lack of data by providing relevant information about the domain, (b) consequently infer new facts about the user, and (c) reinforcing the learning process by introducing newly inferred features.

We start with a brief introduction to neural-symbolic integration model in Section 2. In Section 3, we introduce the building blocks of our proposed approach to build the DMMO ontology as the digital mental model of a user. We continue by providing a short description about the potential domains and applications of the proposed model in Section 4. The paper ends with a brief discussion and conclusion in Section 5.

2. Neural-Symbolic Integration

Neural-symbolic integration has been proposed as a solution to bridge the gap between symbolic and sybsymbolic methods [11, 10], and as a result, to deal with the bottlenecks of one side using the strength of the other side.

Availability of large amount of data has been a strong motivation behind many research focusing on data-driven user modeling by capturing habits and attitudes of a user [12]. Nevertheless, as mentioned above, data-driven and in particular subsymbolic methods suffer from noisy data, sampling bias and lack of good training data [8]. Moreover, even with precise outputs, subsymbolic methods are not able to infer extra information out of the scope of training data. For instance, a recommendation system trained with the data captured from a user’s Spotify account is reliably able to recommend a song in accordance with preferences of the user, however, is ignorant of the user’s personality, despite the possibility to postulate personality traits based on music genre preferences [13, 14].

Neural-symbolic approaches can be employed to tackle the aforementioned issues of data-driven

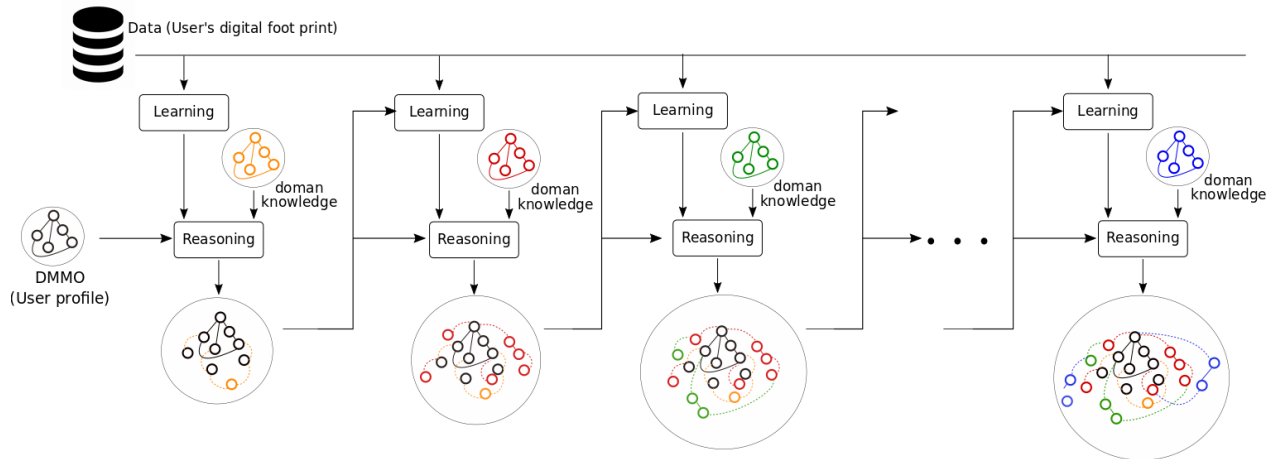


Figure 4: Neural-symbolic integration approach to build an ontological digital mental model of a user (DMMO ontology) based on her digital footprints.

methods by reasoning upon learning outputs and addressing the ambiguities in data [11]. On account of learning competency of their subsymbolic part, they are furthermore able to address the two main issues related to symbolic methods, including (i) the bottleneck of manual knowledge representation, and (ii) poor capability in handling massive data. In the following, we quickly go through the details of symbolic components including knowledge models and reasoning processes.

2.1. Knowledge Models

By a knowledge model we refer to a formal structure that represents a domain in the form of concepts and their relationships, readable both for machines and humans [15]. Providing formally represented human knowledge, knowledge models are indispensable tools to achieve automated reasoning. Many publicly available knowledge models are in the form of ontologies or, in general, knowledge graphs. In recent years, there has been a tremendous increase in the number of interlinked and public ontologies in different areas such as medicine [16], geography [17], music [18], sport [19], etc. There are also a number of upper ontologies, such as DOLCE Ultralite (DUL) [20], that consist of general and domain-independent concepts and relations served as a foundation for domain-specific ontologies [21]. An upper ontology can (indirectly) relate the other domain ontologies already specialized and linked to its general concepts, and thereby, contribute to improving semantic interoperability across multiple domains [22].

2.2. Reasoning

Automated reasoning concerns with enabling machines to perform logic-based reasoning tasks and infer new information out of the known and formally represented knowledge [23]. A computational agent may perform different types of reasoning depending on what it knows (i.e., content of the knowledge model) and what it observes. These types, as listed in Table1, include deductive, inductive and abductive reasoning [24]:

Deductive reasoning is a monotonic reasoning that scrutinizes possibilities to imply a specific logical (and guaranteed) conclusion w.r.t one or a number of general rules or hypotheses given in a knowledge model.

Table 1

Three main types of logic-based reasoning: Deductive, Inductive and Abductive reasoning.

	Deductive Reasoning	Inductive Reasoning	Abductive Reasoning
Given	$A \subseteq B$ (general rule) $a \in A$ (observation)	$a1 \in A, a1 \in B$ (observation) $a2 \in A, a2 \in B$ (observation)	$A \subseteq B$ (general rule) $a \in B$ (observation)
Conclusion	$a \in B$	$A \subseteq B$	$a \in A$
Specification	<ul style="list-style-type: none"> specializing general rules guaranteed conclusion 	<ul style="list-style-type: none"> generalizing observations probably true conclusion 	<ul style="list-style-type: none"> concluding the best guess probably true conclusion

Inductive reasoning refers to an inference process under uncertainty, where a general rule is inferred based on observations. In other words, what the reasoner implies is not guaranteed and can later be modified.

Abductive reasoning is concerned with inferring the best possible explanation (guess) for the given observation. The conclusion of abductive reasoning is not guaranteed either.

2.3. Integration Process

Integration of symbolic methods and learning algorithms has been mainly based on applying symbolic reasoning methods upon the boundary (input or output) layers of a neural network [25]. Such integration, summarized in Figure 5, is able to enhance the transparency of the learning process by explaining its outputs [26]. Apart from transparency, such integration can also improve learning performance and increase its accuracy by closing the loop and sending back the inferred explanation about the output to the input layer as a new channel of data [27]. However, in boundary integration, the reasoner is in fact independent of the structure of the subsymbolic method and is only concerned with the semantics behind the learning outputs.

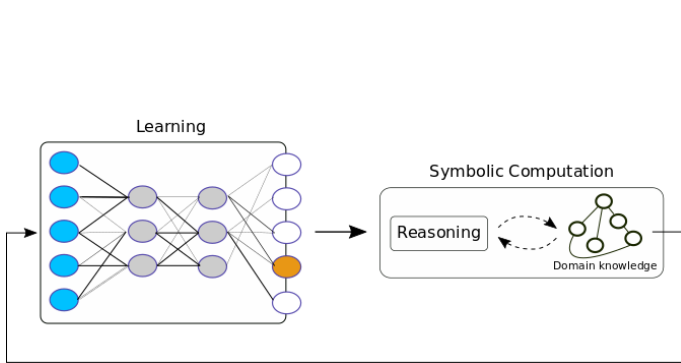


Figure 5: Boundary integration: reasoning is applied upon boundary layers of neural networks.




Figure 6: Hidden layer integration: reasoning is applied upon hidden layers of neural networks.

In addition to boundary integration, instilling reasoning outputs into hidden layers of a neural network has been recently proposed in the literature. The goal behind such entangled approach shown in Figure 6 is to also boost up both the performance and transparency of the learning process, not only by explaining the learning outputs, but also by directly involving the semantics into the learning optimization process [28, 25].

3. Neural-Symbolic User Mental Modeling

A computational mental model (or the digital identity) of a user is a key enabler for having more realistic interactions between humans and robots that comply with the advanced HCI model depicted in Figure 2. To build such a computational model we propose to implement double-loop learning based on an integration of learning and reasoning modules (see Figure 4). In the following sections we go through the main modules of the proposed approach and their role in constructing the user mental model.

3.1. Decentralized Data

The main source of inputs continuously feeding the system is the digital footprints of the user resulted from her interactions with social networks or with different devices such as mobile phones, smart watches, etc. Due to the distributed nature of digital platforms as well as the privacy issues, the generated data may not always be fully accessible. Furthermore, on account of the heterogeneity of data types (e.g., text, image, video, location, environmental data, etc.) perhaps different types of learning methods are required.

To address both the issues of decentralized data and disparate data types we consider federated (or collaborative) learning as a distributed machine learning technique [29]. Federated learning approach enables decentralized platforms holding local data to collaboratively train a model without exchanging data.




Figure 7: An overview of federated learning.

In federated learning, as shown in Figure 7, there is a pool of machine learning models dedicated to each data type (inc. time series, static image or textual format, etc.). Depending on type of data, a suitable model is selected to be trained locally on a platform (e.g., a mobile device, or desktop computer). The locally trained models on separated devices are sent to a server where a federated function is applied to generate and improve the eventual model. Within an iterative process, the improved model is sent back to the local platforms to continue learning from newly recorded data.

Each learning module shown in Figure 4 represents a process executing the federated function on the server and results in a global model updated based on the local learning modules on distributed devices.

3.2. Digital Mental Model Ontology (DMMO):

As shown in Figure 4, the mental modeling generation process is initialized with a knowledge model, more specifically an ontology. This ontology which is henceforth called DMMO (Digital Mental Model Ontology), relies on the upper ontology DOLCE Ultralite (DUL) that represents a fine-grained conceptual framework of general specifications [20]. In the beginning, DMMO holds information about the (source) user profile and is continuously updated in the presence of new data.

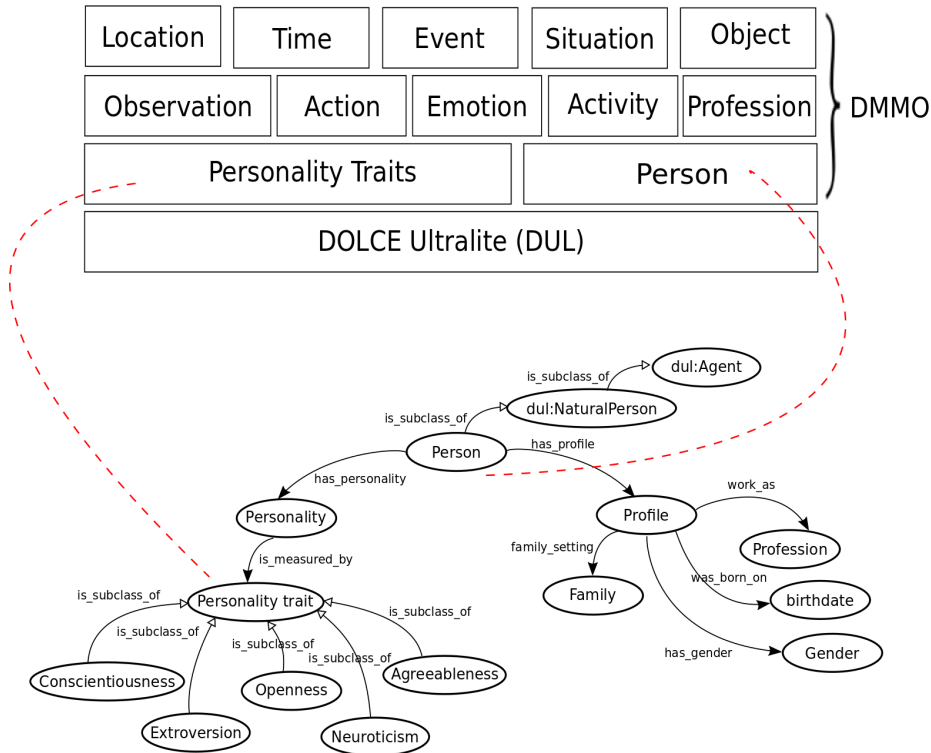


Figure 8: The Digital Mental Model Ontology (DMMO) is a network of ontology patterns representing different aspects of the mental model of a user. The patterns are interlinked through general classes in the upper ontology DOLCE UltraLite. For instance, the two patterns Person and Personality Traits are linked through the class person subsumed by the shared class dul:NaturalPerson.

Figure 8 (top layer) depicts the building blocks of the DMMO ontology where each block will be designed in the form of one or several ontology patterns [30]. Each ontology pattern independently represents a specific aspect of the user mental model such as the general (source) user profile, the daily life of the user, and her personality trait that shows how the person observes (a specific event), acts (in a specific time and location) and feels (about a given situation). The links between the ontology patterns are provided through connections between their superclasses in DUL. Figure 8 (bottom layer) shows how the ontology pattern representing the 5 main personality traits [14] (PersonalityTrait) is connected to the ontology class Person as a subclass of dul:NaturalPerson in DUL, representing the user profile.

3.3. Updating the Knowledge model

Building user digital mental model is an incremental process starting from basic profile of the user, including name, age, gender, profession, family information, etc., represented in DMMO (the Person

ontology pattern). The idea is to populate the DMMO ontology w.r.t the context of the received data and the learning outputs. Figure 9 illustrates the main steps required to build the DMMO ontology as the digital mental model of the (source) user and put it in service to the other (target) users.

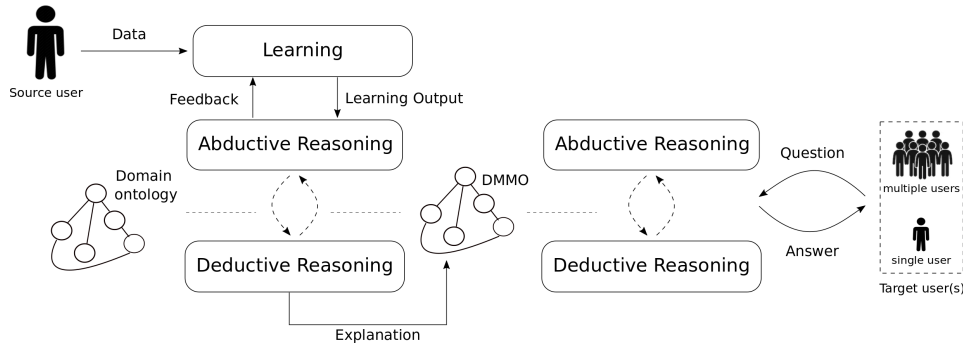


Figure 9: A high level view on building and using DMMO based on interactions between learning and reasoning modules.

To explicate the steps of development, we take an example in the following and describe the technical details related to each module. Suppose that we have gathered data from three social network accounts of a user within a specific period of time (see Figure 10). As mentioned above, the structure of each learning module in our proposed model complies with the federated learning model in Figure 7 which is in charge of the iterative processes to handle decentralized data. However, for the sake of simplicity, we consider a learning module as a pool of different machine learning models used for different purposes such as sentiment classification, activity/scene recognition, time-series forecasting or a combination thereof. Depending on types of the received data, a suitable learning model is chosen and results in learning outputs.

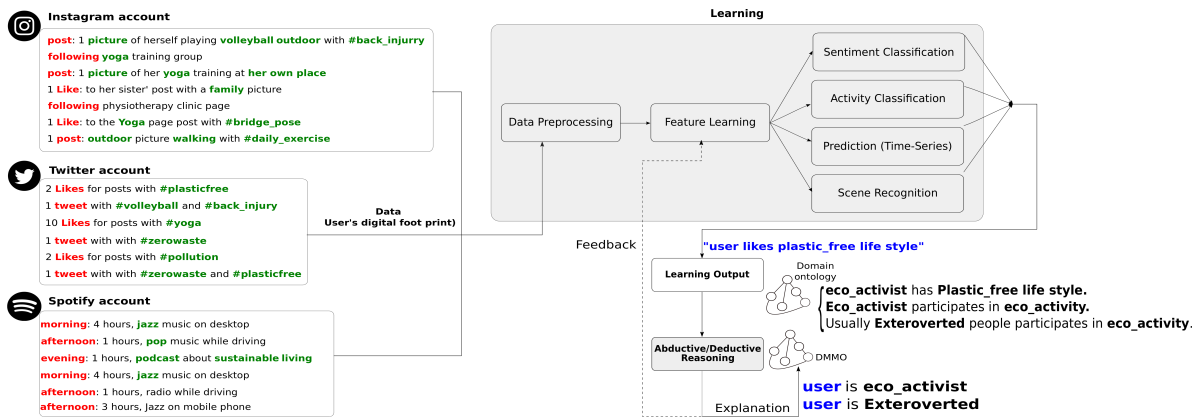


Figure 10: An example showing how the personality of a user is inferred based on her digital footprints.

In the example given in Figure 10, our gathered data is both textual (e.g., captions, hashtags, tweets, etc) and visual (e.g., images). In order to semantically enrich the learning outputs and infer implicit knowledge about the user, each learning process is paired with a reasoning module that is as such associated with the content of DMMO and perhaps the other domain knowledge models. These domain ontologies (e.g., related to sport, music, etc.) are either selected among the existing ones or developed from scratch. Depending on the content of the knowledge model and the features of learning outputs,

different types of reasoning (inc. deductive or abductive) may be applied. For instance, in the given example, after translating the learning output into First Order Logic (FOL) statements, and applying a sequence of abductive and deductive reasoning, we infer the best possible explanation about the personality trait of the user as follows:

Learning outputs:	User likes plastic_free life style.	(1)
Learning outputs into FOL:	likes(user, plastic_free_life_style).	(2)
From Ecology knowledge model:	likes(Eco_activist, plastic_free_life_style).	(3)
Abductive Reasoning on (2) & (3):	Eco_activist(user) .	(4)
From Ecology knowledge model:	participates(Eco_activist, Eco_activity).	(5)
Deductive Reasoning on (4) & (5):	participates(user, Eco_activity) .	(6)
From Psychology knowledge model:	participates(Exterverted, Eco_activity).	(7)
Abductive Reasoning on (6) & (7):	Exterverted(user) .	(8)

As shown in Figure 11, the user is inferred to be extroverted and also high openness because of her music genres of interest. The inferred explanations about the user update the DMMO ontology provided that its consistency is preserved. In case of inconsistencies in the knowledge model, the reasoner requires to verify the learning outputs and their interpretations. After consistency checking, the reasoner sends feedback to the learning module. The feedback includes an explanation based on logical entailments that can identify the causes behind inconsistencies. Feeding the learning process with such information can enhance the performance of the system in better understanding the data, and accordingly in interpreting past events and predicting the future. The model of interaction between learning and reasoning is compatible with the two neural-symbolic integration model explained in subsection 2.3.

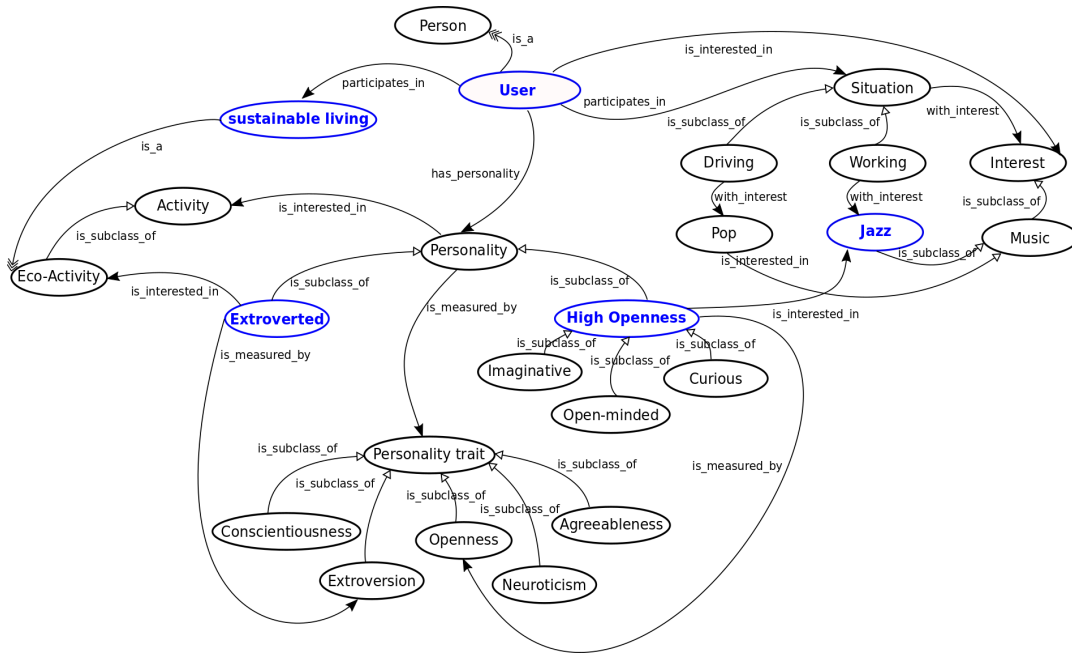


Figure 11: Populated DMMO based on the best possible explanations inferred from data.

However, the DMMO ontology grows incrementally based on the user behaviour captured from data. Moreover, abductive reasoning has non-monotonic nature meaning that the inferred explanations may change in the future. That is why the interaction between learning and reasoning in our

neural-symbolic integration model is bilateral with the possibility to update both sides parameters and achieve an equilibrium.

Equipped with an updated digital mental model of the source user, a computational agent will be capable of interacting with target users and provide answers to their questions w.r.t the personality of the source user. For instance, knowing the source user as an extrovert teacher, the agent which is also equipped with the expert (or professional) knowledge of the user, would probably decide to talk more with its students (target users) during the teaching sessions as extroverts are often more willing to openly talk and clear their opinions [13]. As we can see in Figure 9, the decision of how to respond to the target users is made based on logical entailments as a result of both abductive and deductive reasoning processes upon the knowledge models.

4. Applications

Our proposed framework can result in the emergence of a new application development paradigm that has the ability to combine knowledge representation and data learning in one development framework and makes it easier for a broader group of software engineers to design and build decentralized and more transparent AI-driven applications. Currently we have applied our methodology to the design of four application domains to show the versatility, adaptability and scalability of our framework. The first application is based on our earlier project called “Augmented Eternity” that allows highly-trusted individuals, especially family members, to share affection, knowledge and expertise in the form of digital surrogates and interactive wills. In another application, we demonstrate our methodology in the legal space in which knowledge representation is better structured and we demonstrate how legal experts and lawyers share their expertise in the form of legal avatars, answer general legal questions and we introduce new business models in the legal space. In the third application, we apply our methodology in the education sector, allowing students to wear filters and “lenses” on polarizing subjects to avoid echo chambers measuring the reduction of group-think biases. In the fourth application, we have applied our methodology in a decentralized city recommendation service that allows people to roam cities and wear their trusted relationships as lenses to see the city from their friends point of view rather than relying on centralized travel services like Tripadvisor. Details of these applications and their associated design patterns are presented in our other disseminations. They are all discussed separately in our prospective publications in more software engineering focused venues.

5. Discussion & Conclusion

Our proposed mental model generation approach conforms to the double-loop learning model which is per se implemented based on neural-symbolic integration models. Our motivation behind this hybrid approach is to enable learning methods that instead of solely relying on data-driven methods, they can also employ symbolic reasoning throughout the learning phases in a coherent and complementary fashion. The approach is a step towards modeling human-like learning in which the synergistic interplay of reasoning, analytics, bias and affection recognition is more evident in comparison with software systems. In our proposed approach, the reasoning is applied across the learning process to first recognize mistakes, and then to avoid repeating the similar performed mistakes by adjusting parameters without human interventions during a more unsupervised process. To achieve this target state, we apply symbolic reasoning on hidden layers of neural networks to achieve a more contextual and semantic based inference. We like to mention that the reliability of our proposed approach depends on the domain and category of knowledge models and the ontologies at hand. The com-

plexity of the reasoning methods also depends on how the knowledge is represented in the adopted ontologies (e.g., in terms of the logical operators, quantifiers, and other logical constraints used in the definition of axioms). One of the key hindering factors is the lack of available domain-related causal models that are compatible with the data used for learning. To address these challenges, our approach populates standalone individualistic models about each user under study by leveraging the theories and practices of psychology and cognitive science. These knowledge forms are then mapped as semantic models and will be aligned with existing upper level ontologies on a recurrent and evolutionary basis. Given that a knowledge model holds formalized relations between the parameters of the generic learning model and the features of data, our investigation is required to adjust or extend the available ontologies to satisfy a computationally efficient reasoning process.

Adoption of our approach by other users and researchers can contribute to the formation of a library of ontological patterns that can be plugged in to different domain-driven learning processes such as the financial services, healthcare, retail, legal and entertainment. In future, the software engineering paradigm of these pluggable semantic models can be reviewed in introducing novel design patterns for more robust AI systems. As discussed in the Applications section, we are working with experts in different groups to validate the repeatability and scalability of our approach when the change in domain and semantic models are significant. As part of our future work, we will also be demonstrating how our work enables the privacy-preserved identity of users to become more portable and be used as a standalone software agent for better knowledge and expertise sharing.

References

- [1] A. Calero Valdez, M. Ziefle, K. Verbert, Hci for recommender systems: The past, the present and the future, in: Proceedings of the 10th ACM Conference on Recommender Systems, RecSys '16, Association for Computing Machinery, New York, NY, USA, 2016, p. 123–126.
- [2] N. J. Nersessian, In the theoretician's laboratory: Thought experimenting as mental modeling, *Philosophy of Science Association 2* (1992) 291–301.
- [3] P. N. Johnson-Laird, *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness*, Harvard University Press, USA, 1986.
- [4] D.J. Capelo, C., A feedback learning and mental models perspective on strategic decision making, *Education Tech Research* (2009) 629–644.
- [5] K. Z. Mohmoud, T.M., M. Othman, Enhancing learning process using double-loop theory in purchasing process of healthcare organization toward cutting unnecessary costs, *Journal of Research in Psychology 1* (2019) 1–7.
- [6] J. Beel, B. Gipp, S. Langer, C. Breitingner, Research-paper recommender systems: a literature survey, *International Journal on Digital Libraries 17* (2016) 305–338.
- [7] S. Biundo-Stephan, D. Höller, B. Schattenberg, P. Bercher, Companion-technology: An overview, *KI - Künstliche Intelligenz 30* (2015) 11–20.
- [8] P. Domingos, A few useful things to know about machine learning, *Communications of the ACM 55* (2012) 78–87.
- [9] A. Adadi, M. Berrada, Peeking inside the black-box: A survey on explainable artificial intelligence (xai), *IEEE Access 6* (2018) 52138–52160.
- [10] A. Garcez, M. Gori, L. Lamb, L. Serafini, M. Spranger, S. Tran, Neural-symbolic computing: An effective methodology for principled integration of machine learning and reasoning, *FLAP 6* (2019) 611–632.

- [11] P. Hitzler, F. Bianchi, M. Ebrahimi, M. K. Sarker, Neural-symbolic integration and the semantic web, *Semantic Web* 11 (2020) 3–11.
- [12] J. Wang, Y. Chen, S. Hao, X. Peng, L. Hu, Deep learning for sensor-based activity recognition: A survey, *Pattern Recognition Letters* 119 (2019) 3 – 11. *Deep Learning for Pattern Recognition*.
- [13] F. A. Chamorro-Premuzic T, Personality and music: can traits explain how people use music in everyday life?, in: *Br J Psychol*, volume 98, 2007, pp. 175–185.
- [14] B. Ferwerda, E. Yang, M. Schedl, M. Tkalcic, Personality traits predict music taxonomy preferences, in: *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems, CHI EA '15*, Association for Computing Machinery, New York, NY, USA, 2015, p. 2241–2246.
- [15] R. Brachman, *A structural paradigm for representing knowledge.*, 1978.
- [16] M. A. Musen, *BioPortal*, Springer New York, New York, NY, 2013, pp. 146–147.
- [17] W. V. Siricharoen, U. Pakdeetrakulwong, A survey on ontology-driven geographic information systems, in: *2014 Fourth International Conference on Digital Information and Communication Technology and its Applications (DICTAP)*, 2014, pp. 180–185.
- [18] S. Song, M. Kim, S. Rho, E. Hwang, Music ontology for mood and situation reasoning to support music retrieval and recommendation, in: *Proceedings of the 2009 Third International Conference on Digital Society, ICDS '09*, IEEE Computer Society, USA, 2009, p. 304–309.
- [19] A. Ramkumar, Development of ontology for sports domain, *International Journal for Research in Applied Science and Engineering Technology* (2017) 1244–1248.
- [20] A. Gangemi, N. Guarino, C. Masolo, A. Oltramari, L. Schneider, *Sweetening Ontologies with DOLCE*, Springer, Berlin, Heidelberg, 2002, pp. 166–181.
- [21] V. Mascardi, V. Cordi, P. Rosso, A comparison of upper ontologies, in: *WOA*, 2007.
- [22] L. Elmhadhbi, M.-H. Karray, B. Archimède, Toward the use of upper-level ontologies for semantically interoperable systems: An emergency management use case, in: K. Popplewell, K.-D. Thoben, T. Knothe, R. Poler (Eds.), *Enterprise Interoperability VIII*, Springer International Publishing, Cham, 2019, pp. 131–140.
- [23] J. Harrison, *Handbook of Practical Logic and Automated Reasoning*, 1st ed., Cambridge University Press, USA, 2009.
- [24] A. Conner, L. M. Singletary, R. C. Smith, P. A. Wagner, R. T. Francisco, Identifying kinds of reasoning in collective argumentation, *Mathematical Thinking and Learning* 16 (2014) 181–200.
- [25] M. Gaur, U. Kursuncu, A. Sheth, R. Wickramarachchi, S. Yadav, Knowledge-infused deep learning, in: *Proceedings of the 31st ACM Conference on Hypertext and Social Media, HT '20*, Association for Computing Machinery, New York, NY, USA, 2020, p. 309–310.
- [26] A. Bennetot, J.-L. Laurent, R. Chatila, N. Díaz-Rodríguez, Towards Explainable Neural-Symbolic Visual Reasoning, 2019. Accepted at IJCAI19 Neural-Symbolic Learning and Reasoning Workshop (<https://sites.google.com/view/nesy2019/home>).
- [27] M. Alirezaie, M. Längkvist, M. Sioutis, A. Loutfi, Semantic referee: A neural-symbolic framework for enhancing geospatial semantic segmentation, *Semantic Web* 10 (2019) 863–880.
- [28] M. Alirezaie, M. Längkvist, M. Sioutis, A. Loutfi, Position paper: Reasoning upon learning: A generic neural-symbolic approach, in: *Thirteenth International Workshop on Neural-Symbolic Learning and Reasoning (NeSy)*, 2018.
- [29] K. Bonawitz, H. Eichner, W. Grieskamp, D. Huba, A. Ingerman, V. Ivanov, C. Kiddon, J. Konecny, S. Mazzocchi, H. B. McMahan, T. Van Overveldt, D. Petrou, D. Ramage, J. Roselander, Towards federated learning at scale: System design, 2019. Cite arxiv:1902.01046.
- [30] V. P. A. Gangemi, Ontology design patterns, in: R. S. S. Staab (Ed.), *Handbook of Ontologies*, International Handbooks on Information Systems, 2nd ed., Springer, 2009.