Lifelong Personal Context Recognition[1]

Andrea Bontempelli, Marcelo Rodas Britez, Xiaoyue Li, Haonan Zhao
Luca Erculiani, Stefano Teso, Andrea Passerini, Fausto Giunchiglia
University of Trento, Italy

Abstract. We focus on the development of AIs which live in lifelong symbiosis with a human. The key prerequisite for this task is that the AI understands - at any moment in time - the *personal situational context* that the human is in. We outline the key challenges that this task brings forth, namely *(i)* handling the human-like and ego-centric nature of the the user's context, necessary for understanding and providing useful suggestions, *(ii)* performing lifelong context recognition using machine learning in a way that is robust to change, and *(iii)* maintaining alignment between the AI's and human's representations of the world through continual bidirectional interaction. In this short paper, we summarize our recent attempts at tackling these challenges, discuss the lessons learned, and highlight directions of future research. The main take-away message is that pursuing this project requires research which lies at the intersection of knowledge representation and machine learning. Neither technology can achieve this goal without the other.

Keywords. Personal Situational Context, Knowledge Representation, Machine Learning, Machine-Human alignment

1. Introduction

We focus on the development of AIs which live in lifelong symbiosis with a human and interact with her via smart wearables, for instance, smart phones, smart watches, or also - for certain health-centered applications - medical devices. By *Human-AI Symbiosis* we mean here an AI which is aware of the user's life and well-being, in *all aspects of the user's everyday life*, this being the premise for the AI to provide added value services. It is therefore a *holistic* symbiosis enabling a level of interaction which goes far beyond what achieved so far by personal assistants, which focus on single aspects of the user's life, such as calendar and agenda management, physical activities and fitness, education and learning, information access and management, cf. [1,2,3]. This is a very difficult task whose underlying difficulty, still unsolved, has been known for decades. As discussed in John McCarthy's Turing Award lecture [4], as soon as one gets out of specialized and well defined domains, AIs suffer from the problem of *brittleness*, namely their inability to deal with situations differing even marginally from the scope for which there they were devised. McCarthy called this the *Problem of Generality*. The original formulation was meant only for knowledge-based approaches. For some time it looked like machine

---

learning could provide an effective solution to this problem. However, as discussed in some detail in [5], this turned out not to be the case. As of today, machine learning, and deep learning in particular, are hardly usable in applications which are not restricted in scope or not time-invariant.

The key intuition underlying our approach is that, in a world which is ever changing and which does so in a largely unpredictable way, the only way for an AI to avoid the problem of brittleness is to make sure that its understanding of what is happening is completely aligned with that of its user. In fact, in a world where *(i)* in time, nothing is ever equal to itself and where *(ii)* the meaning of what is happening, i.e., the *intended semantics* of any representation of what is happening, is in the mind of humans and, as a particular case, in the mind of its reference user, humans are an AI's ultimate source of information. In turn, the *only* possible way to carry out this objective is to build an AI that, as a lifelong effort, is able to continually interact with and get relevant information from its reference user, in particular when something new and unforeseen happens. This type of alignment requires full mutual and *bidirectional* explainability between the AI and her reference user. The direction from the user to the machine is what is usually done in Knowledge Representation (KR) and Machine Learning (ML). The other direction has recently emerged under the moniker of Explainable AI (XAI) [6,7] as a major requirement to achieve human-centric and trustworthy AI. The bidirectional aspect of explainability, so far largely unstudied, is however equally and maybe even more important. Via this type of interaction, the AI will become able of acquiring that type of flexibility which allows humans to understand other humans, and also what happens in the real world, even in presence of differing or wrong information.

Based on the above assumptions, the research described here is based on three main ingredients: *(i)* A general KR mechanism for defining the personal context so as to enable the machine to view the world in user's terms. *(ii)* A suite of general ML mechanisms enabling an AI to perform context recognition in the wild, adapting to changes in the world and in its user. *(iii)* A lifelong *machine-human alignment loop* to maintain alignment through bidirectional interaction. The rest of the paper will describe the progress made in the three lines of research mentioned above (Sections 2, 3, 4). Section 5 will report the lessons learned also indicating the implications on the way ahead. Most of this work has been applied in real world scenarios, based on data collected during experiments designed as part of this research. The SmartUnitn2 dataset, built with an experiment involving 158 university students over a period of four weeks (see, e.g., the description in [8]) will be the base for the examples and the results described in the following.

2. Representing the personal situational context

The notion of context used here was originally (informally) defined in [9] as *"a theory of the world which encodes an individual's subjective perspective about it"*. The key intuition underlying the use of contexts is that generality is achieved by moving from the approach where there is only one monolithic theory of the *objective* world to an unbound number of *subjective* views, modeled as contexts, each providing a partial view of the world, i.e., the set of facts which are locally relevant to the current activity [10]. In this work, we represent contexts as Knowledge Graphs (KGs) and assume that the AI is able to store any number of them as (sets of) *Life Sequences* [8], where a Life Sequence is a sequence of contexts. Figure 1 represents a small life sequence with three contexts.
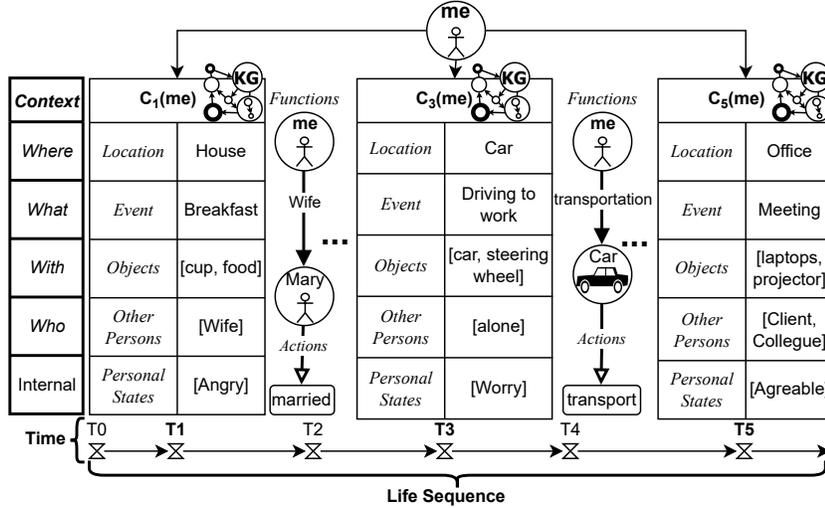
Figure 1. A three context life sequence of me.

Informally, the *personal situational context* describes the circumstances of a person in terms of space-time (Locations and Events), internal context (Personal States), social context (Other Persons), object environment (Objects), and functional relations (Functions and Actions) between the components. Here we talk of *situational* context to mean the context within which a person operates in any moment of time. The situation context is composed by the *Location* of me, which defines the spatial boundaries, and the *Event* within which me is involved at the moment. An event is parameterized on the location as we may have different events occurring in the same location. Location and event are the priors of experience, defining the scenario that needs to be modeled. The change of context coincides with a change in the current location or in the current event.

Given the notion of context as from above, we define the notion of *Life sequence* as a sequence of contexts during a certain period. We assume that me is involved in only one personal context at a time. In fact, at any given moment, a person can be in only one place. This context representation has been used in [11,12].

3. Continually evolving context recognition

Given the notion of personal context, the question becomes how to obtain context information in applications where it is needed. The personal context of a user is typically not directly accessible, at least not in real-time, so it must be inferred from what other information is available. What is needed is a mechanism for enabling an AI to carry out this step. This problem is what we refer to as *personal context recognition* (PCR).

Context recognition in a static world. From one perspective, PCR can be viewed as a generalization of tasks like activity recognition [13] – where one is given access to measurements from handheld or wearable sensors and has to derive what (unobserved) action the user is performing – from a single aspect of the personal context (namely, activity) to the whole context. This immediately suggests an ML approach consisting of two steps: collecting examples of sensor recordings annotated with context information and then using this data to learn a map between the two that generalizes to unseen situa-

tions. A recent investigation of learning techniques for PCR carried out on the SmartU-nitn2 data shows that indeed automated recognition can be achieved with some degree of success [14,15].

Figuring out what predictors and architectures are best suited for this task will likely involve borrowing ideas from activity recognition and related areas, while mixing in strategies for dealing with specific aspects of PCR, such as incrementality and support for interaction, which are critical for lifelong alignment (as discussed below). Another important element is that the personal context is inherently *structured* [15], in the sense that its various aspects are correlated – e.g., a person's activity is strongly influenced by the location that she is in – and constrained by the structure of the context knowledge graph. This hints at the need for developing or repurposing structure-aware predictors. Few or no architectures satisfy all these desiderata and achieve high performance, although progress in this direction is being made [16].

So far, PCR can be viewed as a rather standard (although highly non-trivial) ML task under KR constraints. The real challenges appear when moving to the lifelong setting. Delivering high recognition performance over time requires AIs to be robust to changes *in the world* and *in their user*, which are arbitrary and essentially impossible to anticipate [17], and therefore to be: *(i)* incremental, so as to promptly update as new information is acquired, *(ii)* able to autonomously detect changes and adapt by acquiring appropriate feedback, and *(iii)* able to exploit the prior information stored in the previous contexts, while at the same time, being able to fill them with the knowledge just learned. We unpack these elements next.

The user's description of the context changes. Over time, the user describes the same world in different ways, even if it is not changed. For instance, the user may refer to the same park as central park, urban park, or municipal gardens. At the same time, a relevant fraction of the annotations supplied by users to describe his/her context are unreliable due to mistakes or inattention [18]. Given that acquiring new annotations is expensive and their amount is limited, this issue badly affects the performance of PCR predictors. *Skeptical learning* [18] is a recent interactive learning strategy that tackles annotation inconsistencies. The machine asks the user to revise her annotations if it is confident that there is an inconsistency [18,19]. The machine's uncertainty can be estimated using Bayesian [20] or frequentist [16] techniques, while enabling the AI to explain its skepticism by showing past examples that support the model's suspicion [21]. Both past and new data can be fixed by the user leading to better data and models [21]. This helps the AI to ensure that the information it stores is globally consistent.[2]

The world itself changes. Over time the world changes, and so does the user and her understanding of it. E.g., if the user has to relocate, the structure of her personal context changes. From a statistical perspective, this can be viewed as a form of *concept drift* [23,24] and as such it affects the distribution of sensor observations and of personal context given observations. For instance, during the semester, our user may spend most of her time studying at the library, but once the finals are over, she stops going to the library as often and while there she is less likely to be studying. Here, however, drift is

---

[2]The requirement of consistency was already identified in [4]. McCarthy proposed *non-monotonic reasoning* as a general consistency-preserving mechanism in a changing world. Later, [22] proposed an approach to non-monotonicity based on the use of contexts. However, these and more recent works assume that the user would provide the new information, hindering scalability in practical applications.

more complex as it can affect the knowledge encoded in the context knowledge graph. In other words, whereas concepts like "Friend" and "Library" are essentially immutable, the *specific* friends and libraries that matter to the user do change over time, for instance when she graduates. We refer to this as *knowledge drift* [25]: concepts and relations can become obsolete or irrelevant, and new ones may need to be acquired. Failure to align the AI's understanding to the updated knowledge may lead to providing useless or actively harmful predictions. Given the ego-centric nature of the personal context, the only way to counter knowledge drift is to interact with the user herself. This is actually necessary: different forms of knowledge drift leave a similar footprint on the data, which is thus insufficient to disambiguate between – and therefore properly adapt to – them. In [25], we developed a novel algorithm that tackles all of these issues by integrating *automated* drift detection and adaptation of the machine's knowledge graph with an *interactive* step in which the user helps the machine to disambiguate between alternative kinds of knowledge drift.

4. Machine-human alignment loop

The goal here is to design a general *machine-human alignment loop* that ensures alignment over the life span of the AI. We posit that doing so will involve *continual*, *bidirectional interaction*. The question is then how to structure it such that it is cognitively cheap (on the human's side) and computationally affordable (on the machine's side). This is where most future research lies. For instance, an important issue is how to make sure that the user does not get bothered by a (life)long intensive interaction with the AI (most often not generated by her). A second main issue is that the ML techniques above generate a high number of heterogeneous questions, asking for very different information, motivated by different purposes and based on different background knowledge. Even assuming that the user does not get bothered, how can we make sure that the user does not get confused therefore providing the AI with wrong information?

A second order of problems is that, in order to avoid brittleness, we assume the latter to be based on the context representation in Section 2, an assumption which requires the ability of translating all the possible outcomes of ML in KR. Solving this problem presents a number of difficulties: *(i)* The meaning of different ML labels may be related (e.g., they can be synonyms, or have more/less general meanings, or even be from different natural languages). *(ii)* Labels are polysemous and, as such, have multiple meanings that need to be disambiguated. *(iii)* Labels may be unknown to the AI which therefore needs to extend its vocabulary. *(iv)* The assumption of using labels rather than full text is quite limiting. Some simple versions of this problem have been dealt by the Skeptical Learning [19] and are based on the use of a large multilingual resource, called UKC [26], but we are just at the beginning. The above is at the level of language. But problems exist also at the knowledge level in that a label, even the same label in a different context, may stand for (in the context knowledge graph) for an entity, a class name or property name or value. There is finally a last complexity in that, independent of its meaning, the label learned in a certain moment holds only for that moment, while the local context must represent meaning, and how it is partially stable in time, independently of the local fluctuations. Some very initial work is described in [11].

On the machine side, the need to obtain supervision can be addressed by adapting machine-initiated and human-initiated interaction protocols, such as active learning [27]

guided learning [28]. In real-world lifelong alignment interaction however entails solving additional problems, such as choosing *when* to request annotations so as to maximize response time and quality. Moreover, in order to handle change, the machine will have to interleave requests for supervision with skeptical learning *and* knowledge drift detection and adaptation. How to properly model and implement bidirectional interaction of this kind is left to future work.

5. Lessons learned

What described above are only first steps. Still a few general lessons have already been learned which provide guidelines for the work to come.

The first relates to the KR dimension of the problem. Implementing contexts presents two non trivial difficulties. The first is that the notion of context should allow for the representation of *any* possible real world situation, thus effectively achieving the requirement of generality. The second is that of representing and preserving identity across different contexts. For instance, it must be possible to represent the set of entities, e.g., people, locations, events, which occur across multiple contexts, during the lifetime of a person. Furthermore this must be done in a way that, in such contexts these entities are described by different properties and different, possibly contradicting, property values, this being the key for modeling non monotonic reasoning in a changing world [22].

The second relates to the ML dimension of the problem. Here there is obviously a lot to do in the improvement of the state of the art. But the two issues which seem more pressing and leading to new avenues relate to the full embedding of ML in time (which for instance makes the distinction between training and execution after training meaningless) and to the bi-directionality of the human-AI interaction. It is not a case that the first paragraph in Section 4 does not have citations. This is the exactly the empty space on which we are concentrating now.

The third relates to the Machine-Human alignment loop. The first wave of experiments provided clear evidence that there is a need to provide the AI with Computer Vision (CV) and Natural Language Processing (NLP) capabilities. Only CV allows the AI to construct a model of the world which is similar to that of the human and facilitates their mutual interaction. NLP is needed to scale to the complexity of the world (see Section 4). Here the focus is on building interactions, e.g., in the form of simple dialogues, where the terms used by the AI have the meaning intended by the human. Here the well known and still unsolved *Semantic Gap* problem applies [29]. The work described in [30,31] is a first small step in this direction.

The fourth issue is the importance of running real world experiments in the wild. One reason is that there are no datasets about Human-Machine symbiosis. The second is that there is a need of ethics aware datasets, mainly because of its huge impact on the life of people. The third is that the evaluation of the AI we are developing, which evolves in time while, at the same time, changing the behaviour of the human, requires a careful design of the experiment. Various datasets have been collected, see [32,33,34,35], using the iLog platform [36,8]. A major challenge is the need of a general methodology to be used to run these experiments in a systematic way. This brings up the issue of interdisciplinarity. Initially, the focus was on Philosophy and Cognitive Science, but here there is also a need to learn from the (Computational) Social Sciences (e.g., see [37]).

Author profiles

- Andrea Bontempelli, PhD student, KR, ML and AI
- Marcelo Rodas Britez, Post-doctoral Researcher, KR and AI
- Xiaoyue Li, PhD student, KR and AI
- Haonan Zhao, PhD student, ML and AI
- Luca Eculiani, Post-doctoral Researcher, ML and AI
- Stefano Teso, Assistant Professor, ML and AI
- Andrea Passerini, Associate Professor, ML and AI
- Fausto Giunchiglia, Full Professor, KR, AI and ML

References

[1] Mitchell TM, Caruana R, Freitag D, McDermott J, Zabowski D, et al. Experience with a learning personal assistant. Communications of the ACM. 1994;37(7):80-91.

[2] Guha R, Gupta V, Raghunathan V, Srikant R. User modeling for a personal assistant. In: Proceedings of the Eighth ACM International Conference on Web Search and Data Mining; 2015. p. 275-84.

[3] Dempsey P. The teardown: Google Home personal assistant. Engineering & Technology. 2017;12(3):80-1.

[4] McCarthy J. Generality in artificial intelligence. Communications of the ACM. 1987;30(12):1030-5.

[5] Gini M, Agmon N, Giunchiglia F, Koenig S, Leyton-Brown K. Artificial intelligence in 2027. AI Matters. 2018;4(1):10-20.

[6] Guidotti R, Monreale A, Ruggieri S, Turini F, Giannotti F, Pedreschi D. A Survey of Methods for Explaining Black Box Models. ACM Comput Surv. 2018 aug;51(5).

[7] Miller T. Explanation in artificial intelligence: Insights from the social sciences. Artificial Intelligence. 2018 08;267:1-38.

[8] Giunchiglia F, Bignotti E, Zeni M. Personal context modelling and annotation. In: 2017 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops). IEEE; 2017. p. 117-22.

[9] Giunchiglia F. Contextual reasoning. Epistemologia, special issue on 'I Linguaggi e le Macchine'. 1993;16:345-64.

[10] Bouquet P, Giunchiglia F. Reasoning about theory adequacy. a new solution to the qualification problem. Fundamenta Informaticae. 1995;23(2, 3, 4):247-62.

[11] Giunchiglia F, Britez MR, Bontempelli A, Li X. Streaming and Learning the Personal Context. In: Twelfth International Workshop Modelling and Reasoning in Context; 2021. p. 19.

[12] Xiaoyue L, Marcelo RB, Matteo B, Fausto G. Representing Habits as Streams of Situational Contexts. In: International Conference on Advanced Information Systems Engineering; 2022. .

[13] Chen K, Zhang D, Yao L, Guo B, Yu Z, Liu Y. Deep learning for sensor-based human activity recognition: Overview, challenges, and opportunities. ACM Computing Surveys (CSUR). 2021;54(4):1-40.

[14] Shen Q, Teso S, Zhang W, Xu H, Giunchiglia F. Multi-modal subjective context modelling and recognition. arXiv preprint arXiv:201109671. 2020.

[15] Zhang W, Shen Q, Teso S, Lepri B, Passerini A, Bison I, et al. Putting human behavior predictability in context. EPJ Data Science. 2021;10(1):42.

[16] Teso S, Vergari A. Efficient and Reliable Probabilistic Interactive Learning with Structured Outputs. arXiv preprint arXiv:220208566. 2022.

[17] Dietterich TG. Steps toward robust artificial intelligence. AI Magazine. 2017;38(3):3-24.

[18] Zeni M, Zhang W, Bignotti E, Passerini A, Giunchiglia F. Fixing mislabeling by human annotators leveraging conflict resolution and prior knowledge. Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies. 2019;3(1):1-23.

[19] Zhang W, Zeni M, Passerini A, Giunchiglia F. Skeptical Learning—An Algorithm and a Platform for Dealing with Mislabeling in Personal Context Recognition. Algorithms. 2022;15(4):109.

[20] Bontempelli A, Teso S, Giunchiglia F, Passerini A. Learning in the wild with incremental skeptical Gaussian processes. In: Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence; 2021. p. 2886-92.

[21] Teso S, Bontempelli A, Giunchiglia F, Passerini A. Interactive label cleaning with example-based explanations. Advances in Neural Information Processing Systems. 2021;34.

[22] Giunchiglia F, Weyhrauch RW. A multi-context monotonic axiomatization of inessential non-monotonicity. In: Nardi D, Maes P, editors. Meta-level Architectures and Reflection. North Holland; 1988. p. 271-85.

[23] Tsymbal A. The problem of concept drift: definitions and related work. Computer Science Department, Trinity College Dublin. 2004.

[24] Gama J, Žliobaitė I, Bifet A, Pechenizkiy M, Bouchachia A. A survey on concept drift adaptation. ACM computing surveys (CSUR). 2014;46(4):1-37.

[25] Bontempelli A, Giunchiglia F, Passerini A, Teso S. Human-in-the-loop Handling of Knowledge Drift. arXiv preprint arXiv:210314874. 2021.

[26] Giunchiglia F, Batsuren K, Bella G. Understanding and Exploiting Language Diversity. In: Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-17); 2017. p. 4009-17.

[27] Settles B. Active learning. Synthesis Lectures on Artificial Intelligence and Machine Learning. 2012;6(1):1-114.

[28] Attenberg J, Provost F. Why label when you can search?: alternatives to active learning for applying human resources to build classification models under extreme class imbalance. In: Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining. ACM; 2010. p. 423-32.

[29] Smeulders AW, Worring M, Santini S, Gupta A, Jain R. Content-based image retrieval at the end of the early years. IEEE Transactions on Pattern Analysis & Machine Intelligence. 2000;(12):1349-80.

[30] Erculiani L, Giunchiglia F, Passerini A. Continual egocentric object recognition. arXiv preprint arXiv:191205029. 2019.

[31] Giunchiglia F, Erculiani L, Passerini A. Towards visual semantics. SN Computer Science. 2021;2(6):1-17.

[32] Giunchiglia F, Bison I, Bignotti E, Zeni M, Song D. Trento 2016 - A pilot on the daily routines of University students; 2021. University of Trento Technical Report - DataScientia dataset descriptors. Dataset soon to be available at: `https://ri.internetofus.eu`.

[33] Bison I, Giunchiglia F, Zeni M, Bignotti E, Busso M, Chenu-Abente R. Trento 2018 - An extended pilot on the daily routines of University students; 2021. University of Trento Technical Report - DataScientia dataset descriptors. DataSet soon to be available at `https://ri.internetofus.eu`.

[34] Fausto G, Ivano B, Matteo B, Ronald CA, Marcelo R, Mattia Z, et al.. A worldwide diversity pilot on daily routines and social practices (2020); 2021. University of Trento Technical Report - DataScientia dataset descriptors. Dataset soon to be available at: `https://ri.internetofus.eu`.

[35] Fausto G, Ivano B, Matteo B, Ronald CA, Marcelo R, Mattia Z, et al.. A worldwide diversity pilot on daily routines and social practices (2020-2021); 2022. University of Trento Technical Report - DataScientia dataset descriptors. Dataset soon to be available at: `https://ri.internetofus.eu`.

[36] Zeni M, Zaihrayeu I, Giunchiglia F. Multi-device activity logging. In: Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication; 2014. p. 299-302.

[37] Zeni M, Bison I, Gauckler B, Reis F, Giunchiglia F. Improving time use measurement with personal big data collection - the experience of the European Big Data Hackathon 2019. Journal of Official Statistics. 2020.