TO APPEAR IN:
H. Purohit, V. L. Shalin and A. P. Sheth, "Knowledge Graphs to Empower Humanity-Inspired AI Systems," in IEEE Internet Computing, July-Aug. 2020

Knowledge Graphs to Empower Humanity-inspired AI Systems

Hemant Purohit, George Mason University
Valerie L. Shalin, Wright State University
Amit P. Sheth, University of South Carolina

Abstract

We present a theoretically-motivated design perspective, challenges, and applications of next-generation artificial intelligence (AI) systems. We envision systems with greater capabilities for meaningful human interaction, including socially-adaptive behavior that incorporates personalization and sensitivity to social context and intentionality. Personalized knowledge graphs (KGs) combining generic, common-sense and domain-specific knowledge with both socio-cultural values and norms and individual cognitive models provide a foundation for building *humanity-inspired* AI systems.

Next-generation Intelligent Systems: The Perspective for Designing Humanity-inspired AI Systems

AI appears in a variety of information technology systems that impact daily living (https://bit.ly/30RZeKL), from navigation services to warning proactively about health issues and weather via virtual assistants. Key policymaking and government organizations recognize the daily significance of AI-infused technologies in our society. Indeed, the American AI Initiative based on the presidential executive order (https://bit.ly/303p9jD), the European Union's AI Alliance (https://bit.ly/2CPxIFV), and one of the ten big ideas of the U.S. National Science Foundation (https://bit.ly/3g3M3wL) have all focused on the future of work at the human-technology frontier.

To serve our social needs effectively in different areas of life -- whether healthcare or work performance, AI assistant systems must incorporate and acknowledge the ethical values and social norms that humans take into account when making decisions and taking actions [1]. Consistent with Newell's [2] knowledge level theoretical framework that informs rational behavior of an AI agent, there is a need for principled intelligent systems that reflect and elucidate explicitly both individual and social values and norms in the AI agent's problem solving approach. Humanity-in-the-loop is therefore the core requirement to design such future systems. We define a humanity-inspired AI system as one that incorporates representation and reasoning ability at the knowledge level spanning individual preference and circumstance to collective cultural norms and values. The use of a *humanity-in-the-loop* concept as the core design requirement replaces a worn *human-in-the-loop* conceptualization to account explicitly for the broader socio-cultural perspective in addition to an individual's perspective. Knowledge graphs (KGs) [3] provide the foundation to AI systems for the representation and reasoning capability to address this *humanity-in-the-loop* design requirement.

Key Takeaways:

1. The goal for a humanity-inspired AI system is to achieve the ability to mimic human-to-human interaction, with clear appreciation for and anticipation of personal behavior with respect to context, social norms, and values.
2. The required knowledge of a humanity-inspired AI system is much broader than the past attempts to illustrate Newell's [2] symbol level and knowledge level that emphasize common sense and domain-specific knowledge of the world. The required knowledge level intersects the continuum of knowledge from individual-centric to community-centric values and norms to enable the personalization of reasoning with the continuum of knowledge from common sense to domain-specific knowledge (c.f. Figure 1.)
3. KGs facilitate a natural way for a humanity-inspired AI system to represent and preserve the provenance of continually-acquired knowledge about the behavior and intentionality of interacting humans.
4. A personalized KG within an AI assistant can facilitate the integration of symbolic and statistical computing for personalized reasoning, consistent with a 'top and bottom brain' computational framework [4]. It enables the integration of top-brain driven symbolic computing (empowering AI assistant to attend to value states in the KG for compliance with constraints arising from social norm and values) with bottom-brain driven statistical computing (empowering the AI assistant to learn the distributional representation of value states in the KG and adapting representations under the guidance of perceptions, emergent values and norms).

The recently released 20-year AI research roadmap [5] emphasizes three core research needs, all of which enable the design of such next-generation *humanity-inspired* AI systems: *integrated intelligence, meaningful human-AI interaction, and self-adaptive learning.* Consistent with this strategic research direction, a foundation of KG provides a principled method for AI system designers to represent and incorporate the knowledge at the level required to empower humanity-inspired AI systems in specific application settings.

Below we suggest use-cases that indicate the requirement for a range of knowledge representing values and norms specific to an individual, a group or culture. Our point with these use-cases is not to imply readily-viable applications, but rather to illustrate the breadth of knowledge that would be required for AI systems to make individually and socioculturally meaningful recommendations and enable technology to elevate and enrich the human experience as characterized in [6].

Consider an application concerning elderly healthcare. Virtual assistants based on pure statistical AI computation may not be sufficient to capture the personalized reasoning required for human assistance as we ascend from a lower level (e.g., individual-centric values) to a higher level of relevant knowledge (e.g., family-centric or community-centric values). For example, an abstract value such as *quality of life* is constantly challenged by the interaction of medical guidelines with socializing with family and friends---a general challenge for elder care brought into crisp focus with the prevailing Covid-19 pandemic. The required personalized reasoning for human assistance could be achieved by the integration of statistical AI computation and symbolic AI computation in the AI systems facilitated with personalized KGs.

Figure 1 summarizes the vision to incorporate the different dimensions of high level knowledge for an AI assistant to provide meaningful assistance to humans in varied contexts.

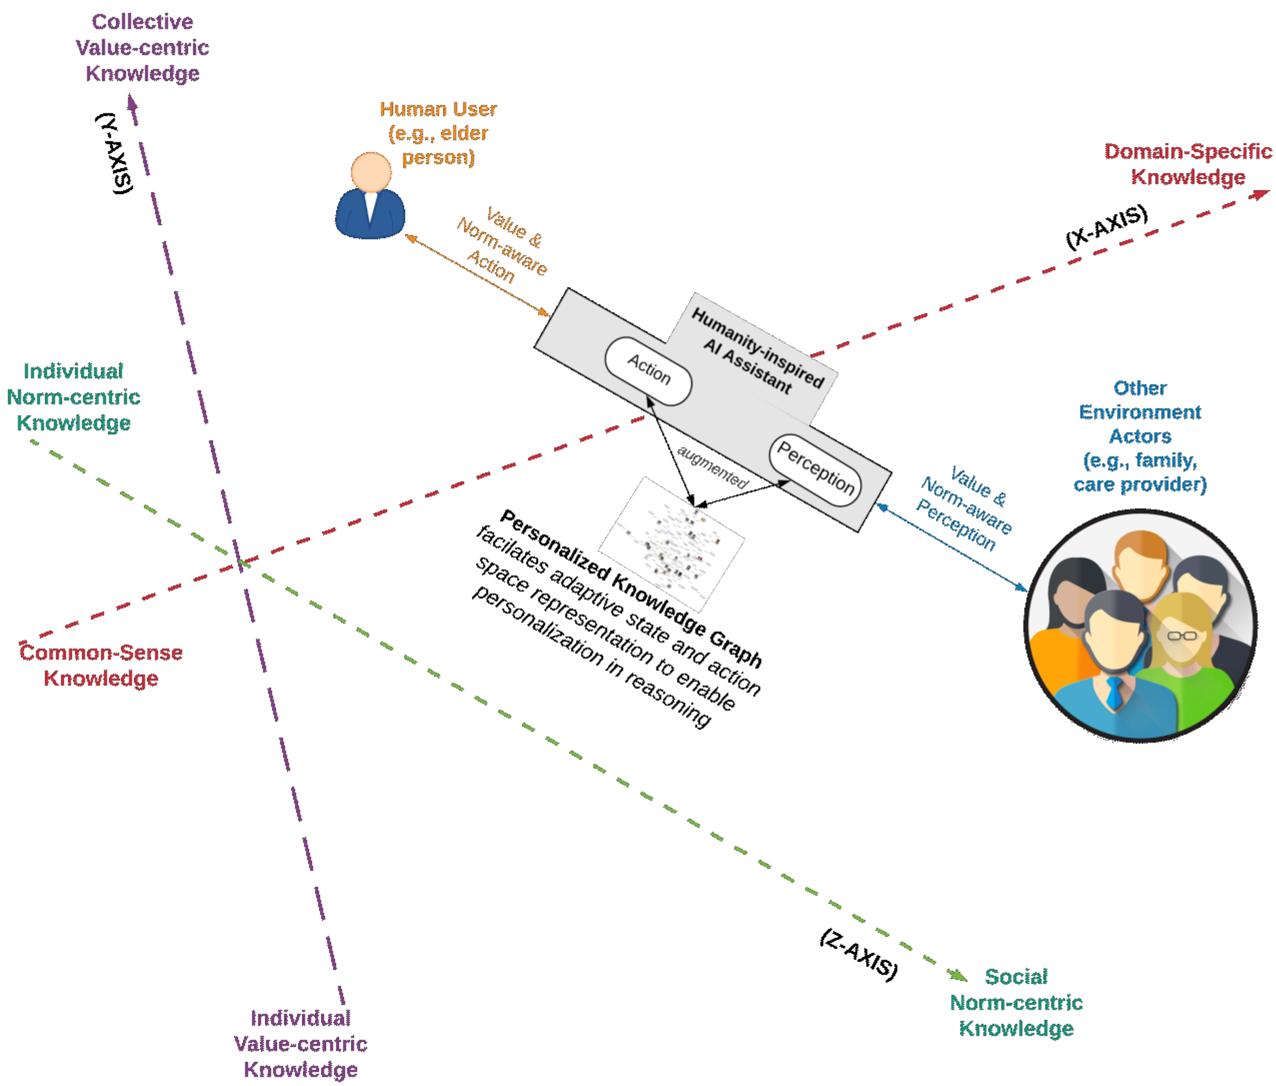

Figure 1. Illustration of a humanity-inspired AI system supported by different dimensions of knowledge. Common-sense to domain-specific context is on the X-axis. Individual to collective values is on the Y-axis. Individual to socio-cultural norm is on the Z-axis. Personalized KGs facilitate contextualized representation and reasoning for such an AI assistant to serve a human with meaningful human-AI interaction capability in different contexts from dyadic interaction to ultimately community-centered engagement.

Illustrative Scenarios of Humanity-inspired AI Systems

We describe four distinct, but related scenarios to illustrate humanity-inspired AI system design and the need to capture multiple dimensions of high-level knowledge to guide the representation and reasoning of an AI system, including not just the continuum of common sense and domain knowledge, but values and norms. The knowledge of values and norms enables the AI system to characterize individual preference and experience within collective groups and cultures.

Consider an elderly person who requires frequent intervention concerning medication and related activities. There are many considerations to provide such care, depending upon the context. We summarize different levels of knowledge required for an AI system to provide such assistance in Table 1 to illustrate the idea of humanity-inspired AI systems.

Table 1. Illustration of different kinds of socio-cultural knowledge required for AI systems to provide human assistance by representing and reasoning across the continuum of values and social norms from an individual to a community of individuals.

Interaction Scenario	Assistance Knowledge Level	Personalization Needs for AI Assistant's State, Action, and Goals	Values & Social Norms
Elder Individual & AI Assistant (e.g., Alexa)	Individual values and common sense knowledge	Scheduling individualized transportation to a medical provider given the preferred value state of *independence*	Person centric
Elder Individual, Nurse, and AI Assistant	Dyadic individual values and norms, and common-sense knowledge	Delaying but not skipping the medication reminder for the elderly patient when outside home, given his/her preference for valuing *privacy* but the care-provider norms of *compliance*	Care centric
Elder Individual, Family, and AI Assistant	Individual and collective values, social norms, and domain knowledge	Preparing for and attending family gatherings and celebrations that disrupt *compliance* with the *care-provider* norms. Context-independent recommendations for meals needed for medication may conflict across the preferences and values of all family stakeholders.	Family centric
Elder Individual, Retirement Community, and AI Assistant	Individual values, social norms and both common-sense and domain knowledge	Scheduling external activities consistent with minimizing risk to *safety* (e.g., from crowds), and maximizing *efficiency* for the elder (e.g., transportation time.)	Community centric

In the first scenario in Table 1, a mobile-based virtual AI assistant like Siri or Alexa can make recommendations to the elderly person. When the elderly person wants to go to a medical provider, he/she may prefer individualized transportation given his/her value for *independence*. Similarly, for reminders to take medication, the AI assistant needs to understand the suitable interaction patterns from prior history as per the timing, when to remind and by using which preferred modality to persuade him/her.

In the second scenario, the AI assistant needs to capture the knowledge of the context of the prescriptions given by the elderly person's doctor as well as the timing and context when to persuade and when to delay the medication reminder to meet the goals of both parties. For instance, the reminder while outside the house could conflict with the preferred value of the elder for *privacy,* yet the medical norms for *compliance* also need to be met.

In the third scenario, the context of the care recipient changes, due to the prioritization of family engagement and socialization. Options for medication delivery now depend on preserving the constraints behind the recommended schedule, time between doses, relationship to meals, and medication interaction. Yet the potential intrusion of context remains limited, and under the control of participants vested in appropriate care.

The fourth scenario illustrates the maximal intrusion of socio-cultural context (and material culture in particular) in care planning, in a context that favors community functionality over the care recipient. For example, planning the attendance of a cultural event must include risk management along with concerns for the medication schedule and constraints. Which option minimizes crowds and travel time? Should a companion carry medication, or will the care recipient return home in time to meet the medication requirements?

All of these illustrative contexts challenge the simplistic notion of alerts and recommendations of currently available AI assistants. Of course, our humanity orientation is not unique to this domain, but equally relevant for examples in emergency response, public health, and environmental policy.

Role of Knowledge Graphs (KGs) for Humanity-inspired AI Systems

KG provides a structured representation for knowledge that is accessible to both humans and machines, with a primary focus on machine comprehension and interoperability of data. As Sheth et al. describes in [3], a KG is often used in a variety of information processing and management tasks such as semantically-enriched applications of search, browsing, recommendation, advertisement, and summarization on the web or socio-technical systems in a variety of domains.

Theoretical Foundation: Knowledge Level Framework.
One of the key foundations of knowledge representation and reasoning in AI systems is the framework proposed by Newell [2] for explaining the rational behavior of an AI agent system. His theory was motivated by three major issues: over-emphasis on specific representations, the controversial focus on theorem proving, and the conflicting opinions on the nature of knowledge. The theory addressed these issues by presenting a vision to differentiate two critical levels for an AI system operation. What the agent knows and its intentions are at the *knowledge level*, in order to rationalize the agent's behavior. On the other hand, how to mechanize the agent's behavior in the system to achieve the desired intent is at the *symbol level*.

Building upon Newell’s knowledge level framework, we present the following three major dimensions of knowledge necessary to design the humanity-inspired AI systems capable of meaningful human interaction (see Figure 1), including socially-adaptive behavior that incorporates personalization and sensitivity to social context and intentionality. Our concern is that these dimensions often remain implicit in the explicit representations and reasoning of legacy AI systems.

a. Common-Sense to Domain-Specific Knowledge

Much of the progress to build applications of AI in general as well as across multiple domains has been driven by the ability to represent knowledge across the continuum of common-sense knowledge to domain knowledge. For instance, knowledge bases like ConceptNet and YAGO have captured and provided the fundamental knowledge of everyday life, leading to many applications of expert systems. Similarly, considering the example of the extensible representation of medical knowledge in the ontologies of Unified Medical Language Systems (UMLS) Metathesaurus, much of the scientific progress at scale has been made possible by the domain knowledge provided in UMLS for all the health science fields including basic biological sciences.

b. Individual Norm-centric to Social Norm-centric Knowledge

Not all physically possible behaviors are acceptable in our daily activities. Thus, an AI system must respect these constraints. The social psychologist Sherif [7] labeled such constraints on behavior “norms”, and much subsequent psychological research has focused on how culture constrains these, especially concerning interaction with artifacts in the material culture such as elevators and sidewalks. In this respect, while preliminary efforts have been made to model norms for interacting agents [8], there is a greater need to realize an environmentally-oriented conceptualization of norms. Anthropologists [9] also note disciplinary, specialized constraints on behavior, conferred for example by training and group membership. Focusing further (with respect to Figure 1) into specialized constraints on behavior, individuals may be governed by personal norms, for example, accommodating limited resources such as time or money. Thus, the primary objective of this dimension of the knowledge level is to ensure the representation of the relevant norms spanning individual circumstance through group membership and culture in the reasoning process of the AI system.

c. Individual Value-centric to Collective Value-centric Knowledge

Rokeach [10] defines values as a class of abstract concepts such as equity, generosity, independence, safety and quality of life that influence individual preferences and choices. Consistent with Rokeach's notion of terminal values, we believe that values determine problematic features of a situation, thereby determining goals and influencing the evaluation of intervention alternatives. Culture confers values, an insight we have exploited in the detection and analysis of radical social media content [11]. However, values are not necessarily coherent within a culture, between individuals within a culture, or within an individual [12]. Furthermore, specific contexts can stress such incoherence of values, for example between independence and safety, forcing individuals to make context-dependent tradeoffs. Thus, the primary objective of this dimension of the knowledge level is to ensure the contextually adaptive representation of individual values in the reasoning process of the AI system.

All of these knowledge dimensions contribute to the effectiveness of exchange between humans, and humans and machines. Although in the early decades of AI medical applications, work from Buchannan and colleagues acknowledged the cultural and individual influences on the design of effective human-machine interaction, their applications were highly constrained. Our vision is much more general, to facilitate a continually-adapted representation of knowledge and contextualized reasoning over a combination of knowledge level dimensions spanning both common-sense and domain knowledge, through culturally and individually determined values and norms, thereby providing a foundation for *humanity-inspired* interactive AI systems. Effective interaction between humans depends on the alignment of divergent, participant-specific situation models informed by all of the dimensions we have identified [13]. KGs <u>must</u> represent these numerous influences on the participants' situation models and can play an important role in the facilitation of communication and joint action when one participant is represented by an adaptive AI system. A KG enables representation, integration, and reasoning across different dimensions of knowledge extracted from diverse modalities and from diverse sources of data in unified form.

Challenges in Building Personalized KGs to Support Humanity-Inspired AI Systems

Realizing the vision of humanity-inspired AI systems will require advancement in both computational and socio-cognitive sciences. We summarize two challenges for each area next.

Integrated Cognitive and Social Science Research

A primary requirement for a humanity-inspired AI system is the ability to apprehend the required knowledge level by incorporating a more well-defined specification of the socio-cultural knowledge for values and norms. Thus, the first major challenge is to explore how to explicitly represent or *learn* to represent the state arguments for values and norms in KGs. A limited representation of prior encoded and learned knowledge may not allow reasoning across different levels of abstraction, in order to both understand the context and make intelligent decisions for recommendation. Cognitive and social scientists can help create resources to learn human-machine interoperable abstractions for the knowledge about values and norms. In this way, the KG construction and updating process can be improved, by not merely providing resources to learn concept category names for values and socio-cultural norms, but rather actively participating in eliciting the in-depth constituent features of the concept category knowledge [14], to guide the learning of statistical representations of concept categories in a KG through their salient features.

Further, the second major challenge is to explore how to guide the reasoning process of the AI system with continually-adapted knowledge in KGs, to incorporate the newly-learned, changing representation of current values and norms in a given context for human assistance. This requires a *parallel* information processing paradigm for an AI system to incorporate an expanded notion of Simon's doing and learning that simultaneously monitors the environment and learns to update its knowledge representation including values and norms, while also meeting the real-time demands of KG search during reasoning.

Computer and Information Sciences Research

In the past decades, the applications of KGs have mainly relied on the generic knowledge bases (e.g., YAGO, DBPedia), with some initiatives for creating KGs specific to a domain (e.g., UMLS Metathesaurus for health sciences). The focus on personalized KGs has not received the attention required to address challenges in both representation and reasoning of AI systems.

The first major challenge is to model incoherent values as concepts and norms as relationship constraints, while creating novel methods to represent a personalized interpretation of concept relationships. Some of these approaches include RDF reification and singleton property [15] based 'statements about statements' in KGs for contextualized and personalized interpretation of knowledge representations in reasoning for humanity-inspired AI systems. For instance, two persons can have the same values (e.g., privacy) but different socio-cultural norms for their behavior (e.g., acceptable personal space depending on culture). Thus, the stored knowledge representation in KGs is very personalized for a human and the AI system should achieve such a level of personalization to perform effective human assistance. To be clear, assumptions about values are built into current recommendation models through user profiling, but their implicit influence on reasoning prohibits explicit personalization.

The second major challenge is rapid learning, and further adaptation, of knowledge representation of values and norms in KGs from few observational data through interaction with the human user. This is analogous to the recent efforts in machine learning research, such as few-shot learning and life-long learning. Although, similar directions need to be further explored to design reliable self-adaptive learning methods for humanity-inspired AI systems to adapt the representations of the knowledge of values and norms over time. The mechanisms [16] to integrate symbolic computing using personalized KGs with representation learning can empower a humanity-inspired AI system, to rapidly learn to attend to those elements in the observational data that correspond to some value states in KGs as well as compliance with social norm constraints.

Conclusion

This article presents a vision for designing humanity-inspired AI systems. Drawing on Newell's Knowledge Level Framework, we use scenarios drawn from elder care to illustrate the essential dimensions of knowledge for an AI system to assist an elder. We advocate for the role of personalized KGs in supporting such AI system designers, and identify challenges for future research and practice.

Acknowledgement

Partial support for underlying research was provided by the National Science Foundation through grants IIS-1815459 and CNS-1513721. Views expressed are those of the authors and not of the sponsors.

References

1. V. Dignum, "Ethics in Artificial Intelligence: Introduction to the Special Issue," Ethics and Information Technology, vol. 20, no. 1, pp. 1-3, 2018, doi: 10.1007/s10676-018-9450-z.
2. A. Newell, "The Knowledge Level," Artificial intelligence, vol. 18, no. 1, pp. 87-127, 1982.
3. A. Sheth, S. Padhee, and A. Gyrard, "Knowledge Graphs and Knowledge Networks: The Story in Brief," IEEE Internet Computing, vol. 23, no. 4, pp. 67-75, 2019, doi: 10.1109/MIC.2019.2928449.
4. A. Sheth, P. Anantharam, and C. Henson, "Semantic, Cognitive, and Perceptual Computing: Paradigms That Shape Human Experience," IEEE Computer, vol. 49, no. 3, pp. 64-72, 2016.
5. Y. Gil, and B. Selman, "A 20-Year Community Roadmap for Artificial Intelligence Research in the US," arXiv preprint arXiv:1908.02624, 2019.
6. A. Sheth, "Computing for the Human Experience: Semantics-Empowered Sensors, Services, and Social Computing on the Ubiquitous Web," IEEE Internet Computing, vol. 14, no. 1, pp. 88-91, 2010.
7. M. Sherif, "The psychology of social norms," New York: Harper, 1936.
8. N. Ajmeri, H. Guo, P. K. Murukannaiah, and M. P. Singh, "Designing ethical personal agents," IEEE Internet Computing, vol. 22, no. 2, pp. 16-22, 2018.
9. E. Livingston, "Making sense of ethnomethodology," London: Routledge & Kegan Paul, 1987.
10. M. Rokeach, "The nature of human values," Free Press, 1973.
11. U. Kursuncu, M. Gaur, C. Castillo, A. Alambo, K. Thirunarayan, V. Shalin, D. Achilov, I. B. Arpinar, and A. Sheth, "Modeling islamist extremist communications on social media using contextual dimensions: Religion, ideology, and hate," Proceedings of the ACM on Human-Computer Interaction, vol. 3, no. CSCW: pp. 1-22, 2019.
12. M. B. Brewer, and S. Roccas, "Individual values, social identity, and optimal distinctiveness," In C. Sedikides & M. B. Brewer (Eds.), Individual self, relational self, collective self (p. 219–237). Psychology Press, 2001.
13. M.J. Pickering, and S. Garrod, "Toward a mechanistic psychology of dialogue," Behavioral and brain sciences, vol. 27, no. 2, pp. 169-90, 2004.
14. M. Yuki, and M. Brewer, "Culture and group processes," Oxford University Press, 2013.
15. V. Nguyen, O. Bodenreider, and A. Sheth, "Don't like RDF reification? making statements about statements using singleton property," In Proceedings of the 23rd international conference on World Wide Web, pp. 759-770. 2014.
16. A. Sheth, M. Gaur, U. Kursuncu, and R. Wickramarachchi. "Shades of knowledge-infused learning for enhancing deep learning." IEEE Internet Computing, vol. 23, no. 6, pp. 54-63, 2019.

Authors:

Hemant Purohit, Ph.D. is an Assistant Professor of Information Sciences & Technology at George Mason University, USA. He can be reached at: http://ist.gmu.edu/~hpurohit

Valerie Shalin, Ph.D. is a Professor of Psychology and Human Factors Area Leader at Wright State University, USA. She can be reached at: https://people.wright.edu/valerie.shalin

Amit Sheth, Ph.D. is the founding director of campus-wide Artificial Intelligence Institute at the University of South Carolina, USA. He is a fellow of IEEE, AAAI and AAAS: http://aiisc.ai/amit