A First Approach Towards Integrating Computational Argumentation in Cognitive Cities

Carlos Iván Chesñevar, María Paula González, Ana Maguitman, Elsa Estevez

Institute of Computer Science and Engineering (ICIC CONICET UNS)
Department of Computer Science and Engineering (DCIC UNS)
CONICET & Universidad Nacional del Sur
Campus Palihue - San Andrés 800
B8000FTN Bahía Blanca, Argentina
Tel. : +54-291-459-5135

E-mail: {cic,mpg,agm,ece}@cs.uns.edu.ar

ABSTRACT

In the last years, the concept of Cognitive Smart City (CSC) emerged from the convergence of the Internet of Things, big data, smart city technologies, and artificial intelligence techniques. At the same time, computational argumentation has consolidated itself as a vibrant area in Artificial Intelligence (AI) which has engineered different approaches to reflect aspects of how humans build, exchange and analyze arguments in their daily lives, mainly to deal with controversial or inconsistent information. Thus, computational argumentation provides a valuable metaphor for reasoning on top of available data in order to draw conclusions and offer explanations for them. This paper discusses a first approach towards integrating computational argumentation in a layered model for cognitive cities, where the bottom layers comprise raw data collected from sensor, actuators and other artifacts deployed in the context of a smart city, and argumentation provides high-level intelligence abilities. We show how our approach can be paired with a layered model for computational argumentation. To illustrate our proposal, we analyze a case study, based on the DECIDE 2.0 framework.

CCS CONCEPTS

Knowledge Representation → Artificial Intelligence → Cognitive Cities → E-government services

KEYWORDS

Explainable Artificial Intelligence, Computational Argumentation, Cognitive Cities Models

1. INTRODUCTION AND MOTIVATION

In the last years, Cognitive Smart Cities (CSC) emerged by the convergence of the Internet of Things, big data, smart city technologies and AI techniques. From the CSC perspective one key challenge is the necessity of novel approaches capable of extracting knowledge from the big data generated from smart cities' sensors [1]. Indeed, not only a new generation of AI strategies are needed but also a human-centered viewpoint with end-user cognition reinforced by cognitive services as a goal. In this setting, one central problem is the requirement of a new generation of explainable AI systems capable of naturally interacting with humans, thus providing comprehensible explanations and awareness of automatically made decisions [2]. Indeed, explainability serves as a means to justify beliefs, reinforce them and enhance their perceived value. Besides, it enables people to diagnose situations, predict or anticipate the future, justify decisions or actions. Explainability also enables people to notice apparent inconsistencies and to restore consistency between their mental models and reality.

Computational argumentation (argumentation) is a critical aspect of decision making. In any argumentation process first, a

ICEGOV2020, April 1-3, 2020, Athens, Greece

ACM ISBN 000-0-0000-0000-0/00/00…$15.00
https://doi.org/00.0000/0000000.0000000

construction of arguments supporting and against a statement is made. Second, the set of warranted (acceptable) arguments is defined, and finally, it is decided whether the statement can be ultimately accepted or not. In the last decade, different computational models were proposed, from pure Argumentation Systems [3] that provide automatic calculus of arguments to more user-oriented tools called Argument Assistant Systems where the goal is to assist the user in the process of arguing, rather than to perform complex reasoning tasks [4]. In between, Hybrid Argumentation Systems (HAS) [5] [6] combine the above two approaches providing software tools for creating, drafting, calculating and analyzing arguments.

This paper discusses a first approach towards integrating computational argumentation in a layered model for cognitive cities, where the bottom layers provide raw data collected from sensors, actuators and other artifacts deployed in the context of a smart city, and argumentation provides the high-level intelligence abilities. Starting from big data collected by city sensors, HAS are proposed as scalable vehicles that support decision making related to rejection or acceptance of premises that underline vast cognitive services responses including awareness of the performed process. By expanding current layered models with computational approaches that emphasize explainability and rational support for decision making, our final goal is to reinforce a citizen-centered viewpoint that enhances awareness, understandable cognitive services and justified services discovery for CSC.

The rest of this paper is structured as follows. Section 2 introduces some background. Section 3 explains the research framework; while Section 4 presents the case study. Finally, Section 5 discusses the preliminary findings and Section 6 summarizes conclusions and future work.

2. BACKGROUND

Smart cities are able to provide services that benefit from the city-scale deployment of sensors, actuators and smart objects [7]. Such services are driven by data, and consequently, they can be broadly classified as producers of data, consumers of data or a combination of both. As pointed out in [1], in a city-scale deployment of smart services data is generated at very high rates, presenting new challenges for smart city designers and developers. Fig. 1 (following the diagram originally presented in [1]) illustrates the challenges for data management of big data (represented by the 3 V's, volume, variety and velocity) and the challenges emerging from machine learning techniques.

As the authors point out [1], in smart city ecosystems there are several features that can be considered from a machine learning perspective: a) humans need to interact with the system to provide their feedback; b) sensors and smart devices generate data at high rate, so that the system should be able to learn and improve from experience; c) a continuous learning mechanism is needed, as smart city applications might evolve over time; d) the data generated by smart city apps is noisy or has some degree of uncertainty.

In [1] the authors point out a number of challenges and future directions, including fast streaming data analytics and context awareness. While smart cities can enormously benefit from machine learning, we contend that this approach alone, in its current state, has an important limitation when providing cognitive services and service discovery: *lack of explainability.* AI has indeed become ubiquitous in society, providing implicit assessment in complex decision making (e.g. movie recommendations in Netflix or tailored advertisement in Google search result pages). However, in many settings in the context of e-government (e.g. e-health medical diagnoses, citizen engagement, tailored smart services, etc.), most AI algorithms suffer from *opacity* (i.e., they are like black boxes - it means that is difficult or impossible to get insight into their internal work mechanism). Consequently, entrusting decisions to such kinds of algorithms raises a number of dangers and ethical questions (see discussion in Section 6). As surveyed in [8], an important amount of scientific research in AI is moving from a "black box" metaphor (associated with several machine learning algorithms such as neural networks and Bayesian belief networks, among others) into a new paradigm called Explainable Artificial Intelligence (XAI) [9]. The ultimate aim is to provide more explainable and transparent AI models while maintaining high-level accuracy.

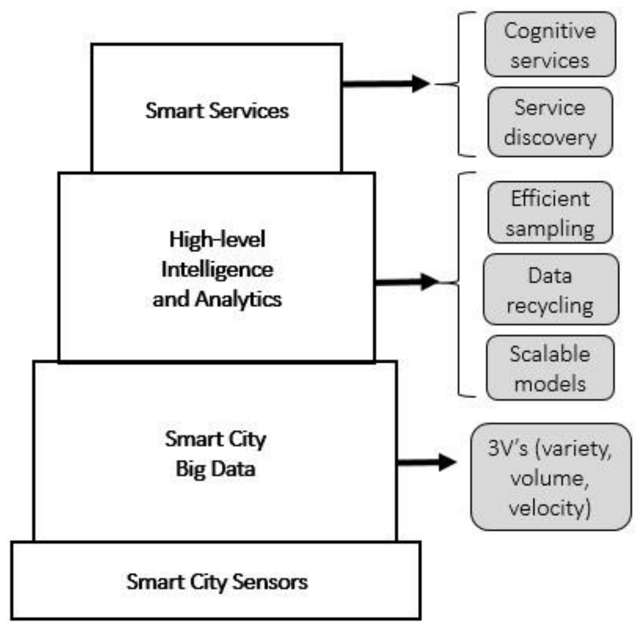

Fig. 1: Challenges of smart cities from a machine learning perspective (adapted from [1])

It must be remarked that XAI has gained momentum in AI from 2016 onwards [8]. On the one hand, the academic community in AI created the so-called FAT* conferences [10] (standing for fairness, accountability and transparency), which bring together AI researchers and practitioners interested in promoting and enabling these features in algorithmic decision making in socio-technical systems. On the other hand, DARPA (Defense Advanced Research Projects Agency, USA) launched its XAI program in 2017, which includes several projects running until 2021 and aims at developing new techniques capable of making intelligent decision systems more explainable [11].

Explainability and transparency are not central concerns of the AI community only. Even though the term CSC is not mentioned explicitly, the United Nation Sustainable Development Goals (SDG) for 2015–2030 includes the urban goal: "inclusive, safe, resilient, and sustainable cities" [12]. In concordance, urban governance faces a wide range of challenges: they need to produce wealth and innovation but also protect the environment and ensure sustainability. Governing a smart city (and consequently a CSC) is about crafting new forms of human collaboration through the use of information and communication technologies [13][30]. In this context, the absence of an underlying formal model makes it hard to provide citizens with a comprehensible and clear explanation of the factors and procedures that led cognitive services to come up with a particular decision or action.

3. RESEARCH FRAMEWORK

In this article, we advocate that computational argumentation offers the proper metaphor for achieving the above goal. As stated in [14], argumentation is an important cognitive capacity that allows us to handle conflicting beliefs, assumptions, viewpoints, opinions, goals and many other kinds of mental attitudes. Argumentation is essential when interacting with other people in a cooperative or competitive fashion to reach a final agreement, or when defending or promoting an individual position.

Computational argumentation has a long-standing tradition in Artificial Intelligence [3]. In general terms, an argument is a tentative (defeasible) proof for reaching a conclusion. Arguments may compete, rebutting each other, so a process of argumentation is a natural result of the search for arguments. Adjudication of competing arguments must be performed, comparing arguments in order to determine what beliefs are ultimately accepted as warranted or justified. Preference among conflicting arguments is defined in terms of a preference criterion. It is difficult to provide a single reference scheme to characterize computational argumentation. However, some layers can be regarded as basic building blocks for the construction of an argumentation model [15]. Five main layers can be identified, as shown in Fig. 2: *structural, relational, dialogical, assessment* and *rhetorical*. Next, we will briefly summarize their main features.

- *Structural layer* - it concerns the structure of arguments and how they are built. In other words, how an argument looks like, in terms of its internal structure. When using a logical language and a knowledge base, an argument can be given by a set of formulae and inference rules leading to a conclusion. When using argumentation schemes (i.e. stereotypical reasoning patterns), a set of premises and a claim are given along with a set of critical questions (i.e., issues that can be raised to challenge arguments built on the basis of a particular schema).

- *Relational layer* - arguments are by nature defeasible, which means that they might be attacked by other better arguments. When considering conflict among arguments, different possible preference criteria emerge (e.g. the most specific argument is preferred over others). The relational layer concerns the interplay in these relationships.

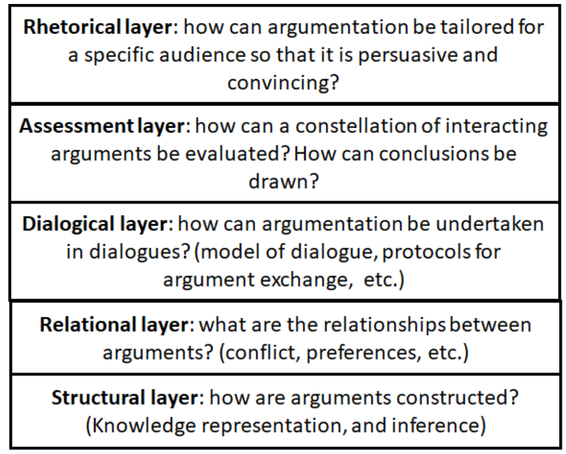

Fig. 2: Key aspects from argumentation (adapted from [15])

- *Dialogical layer* - Arguments can be seen as reasoning pieces that can be exchanged in dialogues (as done by human beings when discussing a particular issue). In the case of argumentation, software agents might engage in the exchange of arguments for a variety of purposes, like inquiry, negotiation, deliberation and persuasion. The dialogical layer characterizes possible protocols to follow for such an exchange.

- *Assessment layer* - once all possible arguments under consideration are on the dialogical arena, acceptability issues come into play. In this stage of the argumentation process, some arguments will be ultimately accepted (so-called warranted arguments), whereas others will be deemed as rejected or undecided. In the end, the outcome of this layer accounts to determine what to believe or what to do, considering all arguments involved.

- *Rhetorical layer* - usually argumentation is undertaken in a specific context of goals. Thus, if the argumentation process aims at persuading some other agent to do something, it is possible to provide rhetorical considerations to address a specific audience or stakeholders. This final layer might be absent in some contexts, or implicitly defined along the argumentation process as a whole.

Our proposal consists in mapping the fourth layer ("High-level intelligence and analytics") proposed in [1] into a "Computational argumentation" layer, where the elements from the previous layers are to be understood as the input from which argumentation and machine learning processes can be carried out. The intended role of this layer is to provide high-level intelligent decision making based on the integration of machine learning and argument-based components, as shown in Fig. 3.

As in the original idea in [1], we consider the same first two layers (sensors, and other artifacts, as a source for collecting big data) along with a "Data Analytics" layer in order to provide an efficient sampling of available information and data reuse. At this point we would like to stress the importance of getting *"labeled* *information items"* as a result, which provide the background for carrying out machine learning and argumentation processes. In fact, a common element along different knowledge representations and knowledge exchange formats is the notion of *label* (e.g. the hashtag in Twitter for marking an information item, the @ sign to identify a user, some special command to identify provenance when dealing with machines in the context of Internet of Things, etc.). Indeed, big data deals with millions of data items that need to be analyzed in order to detect relevant patterns. This accounts for high volumes of low-density, unstructured data, such as Twitter data feeds, clickstreams on a webpage or a mobile app, or real-time sensor information from a particular equipment. Such labels provide the information needed to obtain accurate machine learning models for big data analytics in the context of a smart city's infrastructure.

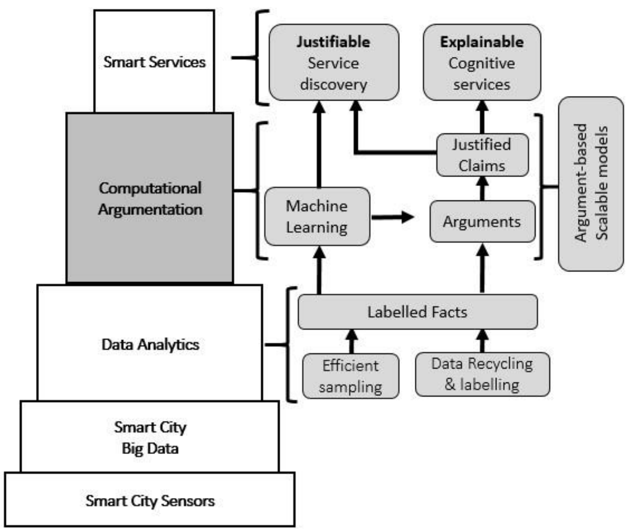

Fig. 3: Proposed architecture for providing smart services based on computational argumentation

Nevertheless, as stated before, we contend that machine learning models alone for big data analytics fall short because of their opacity and limited capacity to enable explainability in AI-powered systems. To this end, we propose that the high-level intelligence obtained from labeled information items should be characterized through machine learning models integrated with computational argumentation. Such integration is naturally a complex issue, but recent research is aiming in this direction (see e.g. [16] [17] [18]). The integration of machine learning models and argumentation should provide a better way to explain underlying reasoning processes, enhancing transparency, accountability and fairness. Decisions from such a high-level intelligence backed by arguments would contribute to harnessing smart services with XAI, aiming at providing a more responsible AI for the different stakeholders involved (decision makers, citizens in general, etc.).

4. RESULTS AND APPLICATIONS. DECIDE 2.0 AS A CASE STUDY

In this section we take a framework developed by the authors called DECIDE 2.0, illustrating how to recast the different components involved in the context of our proposal. Electronic Empowerment Participation captures a novel perspective on e-Participation [19]. In particular, [20] discuss E^2 participation, where collective thinking patterns can be identified under the generic form of "arguments", being contrasted automatically, enhancing thus the abilities of the different stakeholders to engage in creative participatory processes. The underlying machinery that makes DECIDE 2.0 possible is given by agreement technologies [3], a metaphor that integrates several aspects from database theory, artificial intelligence, multi-agent systems and social infrastructures.

DECIDE 2.0 [21] integrates argumentation technologies and context-based search for intelligent processing of citizens' opinion in social media. The framework relies on text mining and opinion mining techniques to filter noise and detect main topics being discussed by citizens in social media. Recognizing that the use of such techniques is not a common government practice, the main contribution of DECIDE 2.0 is to provide an automated tool for extracting arguments based on citizens' opinions. The framework is intended for assessing and confronting pro and con arguments to be used by policy makers and government officials as inputs in decision making. DECIDE 2.0 combines context-based search and argumentation in a collaborative system for managing (retrieving and publishing) service- and policy-related information in social media tools used by governments (for an in-depth analysis of this framework, the reader is referred to [21]).

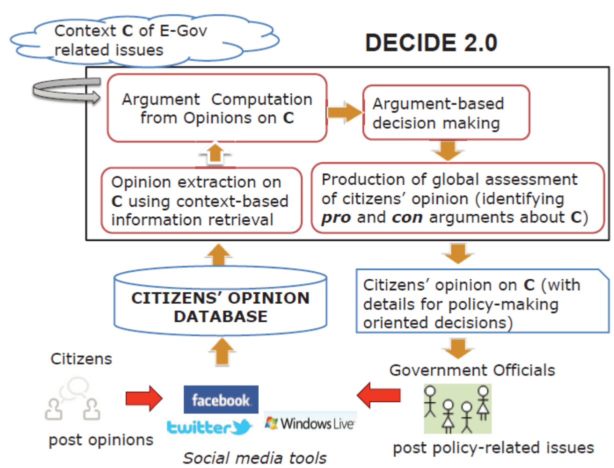

Fig. 4: The DECIDE 2.0 framework

A simplified version of the DECIDE 2.0 architecture is shown in Fig. 4. The architecture comprises different components: a) Opinion Extraction Module –based on data provided by social media, extracts citizens' opinions on a given theme using context-based search, produces formal predicates, and stores opinions and predicates in a knowledge base. It must be remarked that for our analysis we relied on a) Twitter comments posted by citizens (left bottom in Fig. 4) associated with a particular context or topic; b) Argument Computation Module –takes the collected opinions and models them as arguments; c) Argument-based Decision Making Module –based on the generated knowledge base, the component selects predicates on a given theme; d) Global Assessment Generator –based on the results of the previous stage, arguments are classified into pro and con and are consolidated into a global assessment of citizens' opinion on a given theme. The outcome of this assessment is provided to government officials (right bottom in Fig. 4) which can make use of the results to recognize collective thinking patterns. By being able to analyze opinions from mass deliberations of citizens, policy makers can define priorities and better address existing citizenry's concerns.

Fig. 5 illustrates how the different components of the DECIDE 2.0 framework can be mapped onto the different layers in the proposed architecture integrating computational argumentation. It must be remarked that a subset of DECIDE 2.0 was adopted by Mismatica Ltd. and used in the development of a prototype app for city management [22]. In the analysis that follows, we will focus on the model using Twitter information as a basis.

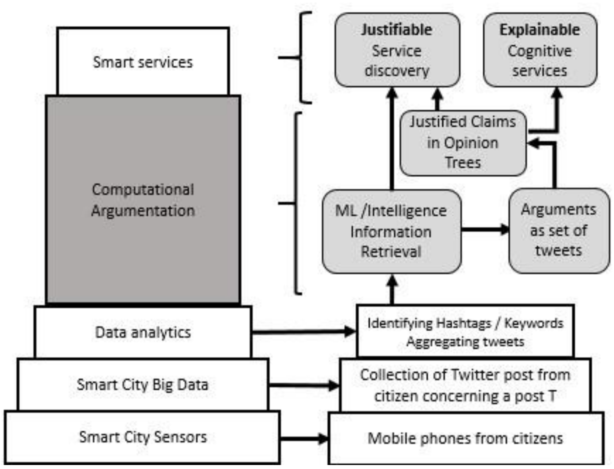

Fig. 5: Conceptualization of the DECIDE 2.0 framework for processing citizen opinion from Twitter under the proposed architecture.

On the bottom layer we have sensors or other artefacts collecting data, e.g. information provided via Twitter from citizens' mobile phones and computers. This input provides the collection of big data which constitutes the citizens' database used in DECIDE 2.0 for performing the argumentation process. The computational argumentation layer, in the case of DECIDE 2.0, involves intelligent information retrieval (in order to collect tweets related to a particular topic or context) to build a model representing collective thinking patterns using sentiment analysis. This model is represented as an "opinion tree", a hierarchical structure that can be incrementally constructed by analyzing the prevailing sentiments (positive, negative or neutral) associated with different aspects of the topic under analysis. Along with that approach, argumentation comes into play as a way of contrasting alternative positions, where more specific information is being used (see Example below). In this particular approach arguments in an opinion tree are just set of tweets referring to a particular topic, whereas the structural relationship among arguments is given by indicating more specific attacks (which are in conflict with a particular argument at issue). The dialogical and assessment layers provide a global view of all the arguments involved in the process, which are finally presented to the government officials.

Consider as an example the following scenario: the traffic management department in a municipality is planning to deliver a project to reconfigure a road junction. The municipality uses its official Twitter account to announce the first stage in this project as follows:

Traffic light controls will be introduced on Stafford Avenue and Goldsmith Street in an effort to improve road safety #StaffordGoldsmith

The announcement prompted immediate and mixed reactions from the citizens. While some welcome the initiative others argue against it. Some examples of tweets posted by the citizens in response to the announcement are the following:

When divers see a green light at #StaffordGoldsmith, will they be watching the road? They will probably be watching the light when your grandmother or son crosses the street 😕 *Commonsense is better than traffic lights* 👎

Expect very long delays soon at #StaffordGoldsmith, especially if you are going towards downtown during rush hours

👏👏👏*Another good initiative from our municipality. Soon we will be able to drive more safely #StaffordGoldsmith*

Are you planning to stop on the red light at 3am in this unsafe neighborhood??? I'd rather be ticketed than robbed!!! #StaffordGoldsmith

The traffic light control at #StaffordGoldsmith is ok as long as it is a smart one!

An analysis based on sentiment analysis on the set of tweets containing the hashtag #StaffordGoldsmith gives as a result the opinion tree shown in Fig. 6. The algorithmic procedure for building an opinion tree based on a collection of tweets is described in detail in [23].

The resulting opinion tree sheds light on collective thinking patterns regarding the announced initiative. It is possible to see that the general opinion about this project is mixed (marked with "=" on the root node of the opinion tree). However, when different aspects associated with the initiative are analyzed, different stands arise in favor and against the project. In particular, several tweets refer to the fact that traffic lights will favor safety (marked as "+" on the nodes labeled *safety* and *drivers*). However, others argue that traffic lights are not necessarily safe for pedestrians, unless the new traffic light control system is equipped with pedestrian traffic lights (marked as "-" on the node labeled *pedestrians)*. Others will argue that although traffic lights tend to improve drivers' safety, it won't be safe to stop at that particular junction late at night because the neighborhood is not safe (marked as "-" on the node labeled *unsafe neighborhood*). Another set of tweets claim that rather than improving traffic flow, traffic lights produce delays and more congestion (marked as "-" on the node labeled *delays*), but some also mention that delays could be significantly reduced if a smart traffic light control system is implemented (marked as "+" on the node marked as *smart traffic lights*). The discussion could go further, including claims such as the fact that idling cars are a major source of greenhouse emission and other statements in favor or against the initiative.

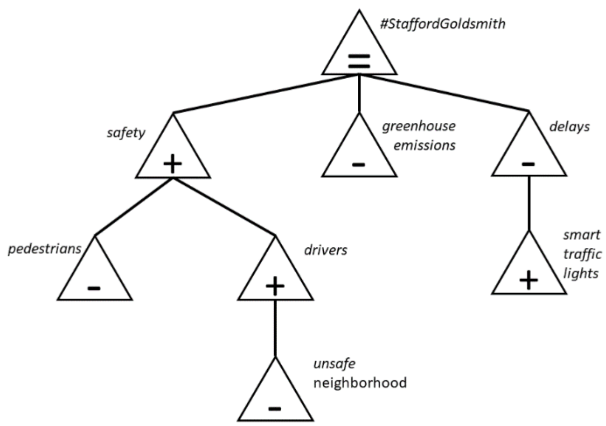

Figure 6: Opinion tree associated with the announcement made by the traffic management department concerning the implementation of a traffic light control system at the Stafford Avenue and Goldsmith Street junction.

By producing a structured global assessment of citizens' opinions in the form of an opinion tree, the DECIDE 2.0 system empowers the traffic management department to reach a more informed decision regarding the initiative. After analyzing the collective arguments in favor of and against the traffic lights proposal, they decide to advance with the project with the following provisions:

- The traffic lights control system will be equipped with pedestrian traffic lights.
- The system will be outfitted with cameras, motion sensors and artificial intelligence software to reduce vehicle idling time, thus avoiding unnecessary delays and greenhouse gas emission.
- The traffic lights will switch to flashing yellow between 10pm and 6am to reduce risks for drivers of being victims of crime.

5. RELATED WORK AND DISCUSSION

As mentioned in [11], there is an urgent need to find ways to explain AI system's predictions to the decision makers so that they know that their decisions are reasonable (simply invoking a neurological metaphor might not be sufficient). The goals of explanation involve persuasion, but that comes only as a consequence of understanding how the AI system works, the mistakes the system can make, and the safety measures for it. Indeed, governments and citizens are expressing concern about the emerging "black box society" [24]. New laws –as the European Union's General Data Protection Regulation (EU) 2016/679 (GDPR) regarding citizen data protection and privacy– prohibit "automatic processing" unless user's rights are safeguarded. Users have a "right to an explanation" concerning algorithm-created decisions that are based on personal information, thus restricting and banning a considerable set of AI systems [25].

In accordance, [26] presents an up to date review about explanations in AI. Authors distinguish between machine-interpretable models and human-interpretable models, as well as between explanation and justification. Various kinds of AI systems were included, and two main branches of current research were differentiated: interpretable models, and prediction interpretation and justification. In [27] it is argued that interpretability is multidimensional. Some possible dimensions include: explainability, accountability, fairness, transparency, and functionality. In particular, [27] considers as a key challenge the need to rationalize, justify, and understand systems and users' confidence in their results.

To cope with the above scenario, this paper proposes a characterization of CSC cognitive services through computational argumentation. As can be seen in Fig. 3, the novel approach aims at providing citizens and policy makers with graphical argument-based information that justifies claims. In this way, the evolution of current smart services into cognitive services that grant awareness and transparency about the decision making process they perform is enhanced. Besides, collective argumentation issues are considered, as it is possible to go beyond the individual interaction of a particular citizen with a cognitive service [28]. Indeed, by means of the sampling and data recycling-labeling process, the capability of the Computational Argumentation layout of Fig. 3 not only provides a solution to a single end-user but also allows to offer explainable cognitive services to a target audience. As an example, the conceptualization of the DECIDE 2.0 framework (see Fig. 5) describes how to treat collective citizen opinions related to a topic.

6. CONCLUSIONS AND FUTURE WORK

CSC are clearly paving the way for a new conceptualization of smart cities, in which computational intelligence, big data, and several technologies converge to offer smart services and tools. As we have discussed in this paper, machine learning alone does not provide the proper insight for conceptualizing many aspects involved in intelligent decision making (particularly those related to explainability issues). In this respect, computational argumentation proves to be a powerful complement for combining the best of both worlds, allowing to identify different perspectives based upon big data from different sensors, while at the same time providing a rational procedure for contrasting alternative perspectives in which there might be potentially inconsistent or incomplete information. We have outlined an approach for integrating argumentation into the layered model for CSC proposed in [1], and shown how a particular case study (the DECIDE 2.0 framework) can be seen as a particular instance of this approach.

Part of our future work is focused on enhancing a full-fledged framework for integrating computational argumentation in CSC. This involves considering specific issues related to context-based reasoning and preference among arguments (i.e. different contexts give rise to different kinds of argumentation, using different preference orders). Besides, the problem of the *knowledge acquisition bottleneck* described in [15] should be considered. Indeed, automatic translations are needed between forms in which the citizen knowledge is usually expressed, e.g., natural language, and forms that can be automatically processed by computational argumentation, especially in the smart services layout of our proposal. In particular, complementing traditional visualization techniques for argument-based computation (usually tree-shaped schemes such as the one shown in Fig. 6) with their associated explanations in natural language would be a major advance to enhance explainability. Some alternatives as the one reported in [29] are being explored.

Another key challenge for the applicability of our layered model is the evaluation of the justified claims (see Figure 3). Indeed, suppose the goal is to have a human-centered viewpoint with end-user cognition reinforced by cognitive services. Looking for explainability implies further justifying claims, by addressing factors such as understandability and affordance, consistency with citizens' reasoning and intuition, and expressivity and adequacy, among others. It seems that this cannot simply be achieved by adding a series of psychological experiments or traditional usability oriented tasks at the top layer, as the collective nature of CSC goes beyond current User Experience considerations. Indeed, a revision of what we consider individual and collaborative human-data interaction in CSC scenarios must be carried out, including the analysis of their characterization and relevance in the layered architecture presented here. Research work in this direction is currently underway.

ACKNOWLEDGMENTS

This work was supported by Universidad Nacional del Sur (PGI-UNS 24/N045 y 24/N051) and ANPCyT (PICT 2014-0624).

REFERENCES

[1] Mehidi Mohammadi and Ala Al-Fuqaha, 2018. Enabling Cognitive Smart Cities Using Big Data and Machine Learning: Approaches and Challenges. IEEE Comm. Magazine, pp. 94–101. DOI: 10.1109/MCOM.2018.1700298

[2] Jose M. Alonso and Corrado Mencar. 2017. Building Cognitive Cities with Explainable Artificial Intelligent Systems. Proc. of the 1st Intl. Workshop on Comprehensibility and Explanation in AI and ML (CEx@AI*IA), Bari, Italy, November 16-17.

[3] Sanjay Modgil et al. 2012. The Added Value of Argumentation. In Agreement Technology Handbook, vol. Nro 8, 357—404. In the Law, Governance and Technology Series. Sascha Ossowsky (Ed) Springer Verlag.

[4] Jörg Pührer. 2017. ArgueApply: A Mobile App for Argumentation. In: Balduccini M., Janhunen T. (eds) Logic Programming and Nonmonotonic Reasoning. LNCS, vol 10377. Springer. DOI 0.1007/978-3-319-61660-5_23

[5] María P. González, Sebastián Gottifredi Alejandro García and Guillermo Simari. 2011. Towards Argument Representational Tools for Hybrid Argumentation Systems). LNCS, vol 6772, Springer-Verlag, ISSN 0302-9743, pp. 236-245. USA.

[6] Lucas Carstens, Xiuyi Fan, Yang Gao and Francesca Toni. 2015. An Overview of Argumentation Frameworks for Decision Support. In: Croitoru M., Marquis P., Rudolph S., Stapleton G. (eds) Graph Structures for Knowledge Representation and Reasoning. GKR 2015. LNCS, vol 9501. Springer. DOI 10.1007/978-3-319-28702-7_3

[7] Ala Al-Fuqaha , Mohsen Guizani , Mehdi Mohammadi, Mohammed Aledhari and Moussa Ayyash. 2015. Internet of Things: A Survey on Enabling Technologies, Protocols and Applications. IEEE Comm. Surveys & Tutorials, vol. 17, no.4, pp.2347-2376. DOI: 10.1109/COMST.2015.2444095

[8] Amina Adadi and Mohammed Berrada. 2018. Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI). IEEE Access, Vol. 6, p. 52138-52160, Feb. 2018. DOI: 10.1109/ACCESS.2018.2870052

[9] Shane Mueller, Robert R. Hoffman, William Clancey, Abigail Emrey and Gary Klein. 2019. Explanations in Human AI Systems: A Literature Met-Review. Synopsis of Key Ideas and Publications and Bibliography for Explainable AI. CoRR, arXiv:1902.01876v1 Retrieved from https://arxiv.org/abs/1902.01876

[10] ACM Conference on Fairness, Accountability, and Transparency. See https://fatconference.org

[11] David Gunning and David Aha. 2019. DARPA's Explainable Artificial Intelligence (XAI) Program. AI Magazine, v 40 (2), 44-58. Association for the Advancement of Artificial Intelligence. ISSN 0738-4602

[12] United Nations. The Sustainable Development Goals 2015-2030. Available at https://una-gp.org/the-sustainable-development-goals-2015-2030/

[13] Albert Meijer and Pedro Rodriguez Bolivar. 2016. Governing the smart city: a review of the literature on smart urban governance. International Review of Administrative Sciences (IRAS), vol 82 issue: 2, page(s): 392-408 ISBN 00208523. DOI https://doi.org/10.1177/0020852314564308

[14] Katie Atkinson, Pietro Baroni, Massimiliano Giacomin, Anthony Hunter, Henry Prakken, Chris Reed, Guillermo Simari, Matthias Thimm and Serena Villata. 2017. Towards Artificial Argumentation. AI Magazine. 38(3), 25-36. DOI https://doi.org/10.1609/aimag.v38i3.2704

[15] Douglas Walton, Katie Atkinson, Trevor Bench-Capon, Adam Wyner and Dan Cartwright. 2010. Argumentation in the framework of deliberation dialogue. In Arguing Global Governance: Agency, Lifeworld and Shared Reasoning, C. Bjola and M. Kornprobst, eds, Routledge, pp. 210–230.

[16] Martin Mozina, Jure Zabkar and Ian Bratko. 2007. Argument based machine learning. In Argumentation in AI. Bench-Capon, T, Dunne, P (eds). Artificial Intelligence, v171, issue 10-15,pp 922-937

[17] Guillermo Simari and Iyad Rawan (Eds.). 2009. Argumentation in Artificial Intelligence. Springer Verlag, ISBN 978-0-387-98197-0. DOI https://doi.org/10.1007/978-0-387-98197-0

[18] Fuyuki Ishikawa and Yutaka Matsuno. 2018. Continuous Argument Engineering: Tackling Uncertainty in Machine Learning Based Systems. In: Gallina B., Skavhaug A., Schoitsch E., Bitsch F. (eds) Computer Safety, Reliability, and Security. SAFECOMP 2018. Lecture Notes in Computer Science, vol 11094. Springer. DOI https://doi.org/10.1007/978-3-319-99229-7_2

[19] Sourav Mukherjee. 2019. How IT allows E-Participation in Policy-Making Process. In CoRR. arXiv:1903.00831v2 Retrieved from https://arxiv.org/abs/1903.00831

[20] Carlos Chesñevar, Ana Maguitman, Elsa Estevez, Nardine Osman and Carles Sierra. 2013. E2 participation: electronically empowering citizens for social innovation through agreement technologies. In 14th Annual International Conference on Digital Government Research.

Mellouli, S, Luna-Reyes, L, Zhang, J (eds). ACM, ISBN 978-1-4503-2057-3.

[21] Elsa Estevez, Carlos Chesñevar, Ana Maguitman, and Ramón Brena. 2012. DECIDE 2.0: a framework for intelligent processing of citizens' opinion in social media. In Proceedings of the 13th Annual International Conference on Digital Government Research (pp. 266-267). ACM.

[22] Florencia Gamboa, Marcos Malamud, Ana Maguitman and Carlos Chesñevar. 2016. Citymis optree: Intelligent citizen management using sentiment analysis and opinion trees. In Proceedings of the 9th International Conference on Theory and Practice of Electronic Governance (pp. 250-253). ACM.

[23] Kathrin Grosse, María P. González, Carlos Chesñevar and Ana Maguitman. 2015. Integrating argumentation and sentiment analysis for mining opinions from Twitter. AI Communications, 28(3), 387-401.

[24] Frank Pascuale. 2015. The Black Box Society. Harvard University Press. ISBN 0674368274.

[25] Bryce Goodman and Seth Flaxman. 2016. EU regulations on algorithmic decision-making and a “right to explanation”. Proc. of the 2016 ICML Workshop on Human Interpretability in Machine Learning (WHI 2016). arXiv:1606.08813v3 Retrieved from https://arxiv.org/abs/1606.08813

[26] Or Biran and Courtenay Cotton (Eds.). 2017. Explanation and Justification in Machine Learning: A Survey. IJCAI 2017 Workshop on Explainable Artificial Intelligence (XAI).

[27] Supriyo Chakraborty et al. 2017. Interpretability of deep learning models: A survey of results. 1-6. 10.1109/UIC-ATC.2017.8397411.

[28] Gustavo Bodanza, Fernando Tohmé and Marcelo Auday. 2017. Collective argumentation: A Survey of Aggregation Issues around argumentation. frameworks. Argument & Computation, vol. 8, no. 1, pp. 1-34 DOI: 10.3233/AAC-160014

[29] Theodoros Mitsikas, Nikolaos Spanoudakis, Petros Stefaneas and Antonis Kakas. 2017. From Natural Language to Argumentation and Cognitive Systems. In Proc 13th International Symposium on Commonsense Reasoning (Commonsense-2017), University College London, UK,

[30] Milá Gasco-Hernández. 2018. Building a smart city: lessons from Barcelona. Communications of the ACM, 61(4), pp. 50-57, ACM, USA.