Natural Language Generation Requirements for Social Robots in Sub-Saharan Africa

C. Maria KEET[1]
[1] *University of Cape Town, 18 University Avenue, Cape Town, 7701, South Africa*
Tel: +27 21 650 2667, Fax: + 27 21 650 3551, Email: mkeet@cs.uct.ac.za

Abstract: Robots are deployed in Africa mainly in manufacturing, yet they may assist in society as future oriented technologies as well. They may ameliorate, e.g., service delivery issues and skills shortages. In this discussion paper, several uses and use cases relevant to Sub-Saharan Africa are described and requirements identified. We zoom in on human-robot interaction in Niger-Congo B ('bantu') languages. Use cases for healthcare and education elucidate specific requirements for the natural language generation component of robots in society. In contrast to typical generation systems, it demands for i) combining data-to-text and knowledge-to-text in one system, ii) generating different types of sentences so as to switch between written and spoken language, and iii) processing non-trivial numbers.

Keywords: Robots, Niger-Congo B (Bantu) Languages, Natural Language Generation, Technology-enabled Healthcare and Learning.

1. Introduction

Artificial Intelligence receives ample attention in Africa, both in the broad sense and particular subtopics, such as deep learning applications and the 4th industrial revolution[i]. The latter in particular has as component *robots* (e.g., [4]), be it as movable or unmovable devices or software-based to do certain tasks automatically. Robots are both developed in Africa[ii] and imported, in particular for society-oriented robots, such as digital assistants like the Amazon Echo. The imported robots answer with an American-English voice, rather than a local accent or any of the indigenous languages. Moreover, social robots require content interaction, yet the answers they provide are heavily influenced by the "Big Tech" corporations and an international user-base in the Global North. Such question-answering algorithms use online content, of which there is relatively little on Africa. In addition, a question's answer is based on global priority, since, unlike Web-based searches, Digital Assistants prohibit manual configuration of their search scope; e.g., asking Alexa in South Africa "What is the EFF?", it responds with "Electronic Frontier Foundation", rather than "Economic Freedom Fighters" that a South African would expect.

Thus, when a robot has to interact with humans, it is not yet adapted to the context. The first step toward problem specification, then, concerns the answer to: What does it mean to have them be *culturally aware* and *suitable or appropriate for societies in Africa*? Zooming in on one aspect of it specifically, being language as part of one's identity: where could robots be used to assist, rather than try to replace, humans, and what implications does that have for the language technologies? Within that scope more specifically and given space limitations of this paper: what is the state of the art and what are the unmet requirements for generating personalised interaction, and generation of text specifically?

In this brief discussion paper on future-oriented technologies, we analyse the landscape on human-robot interaction (HRI) with a focus on a language needs analysis for African languages, and in particular the natural language generation (NLG) features and requirements for contextually relevant social robots. While some disparate advances have been made, the slightly different application scenarios of the use cases revealed a substantial set of requirements that will have to be met for successful deployment.

In the remainder of this paper, we first outline the objectives and methodology in Section 2. We contextualise robots for Africa in Section 3 and proceed to natural language generation requirements and illustrative use cases in Section 4. We close in Section 5.

2. Objectives and Methodology

The main objective of this paper is to discuss robots as future-oriented technologies with use cases that focus the language they use for HRI, where the language of interaction belongs to the Niger-Congo B (NCB) family of languages (aka Bantu languages), which are still widely spoken in social settings and seen as part of one's identity. We narrow down this topic and the methodological approach to: (1) Identify scenarios where robots may assist activities in society, as exploratory analysis for robots in the light of ameliorating societal challenges; (2) Determine the mode of language-oriented HRI required for the scenarios; (3) Considering two small use cases, provide an overview of current capabilities for text-based interactions for languages in the NCB family of languages, and the automatic generation of responses in particular; and (4) Assess shortcomings and possible future research directions for response generation in HRI for NCB languages.

3. Technology: Contextualising Robots

There are different definitions of robot, but they agree on that it is a machine that carries out tasks or actions automatically that may or may not be human-like actions. A robot does not need to be humanoid and could include robot vacuum cleaners, autonomous vehicles ('robot chauffeurs'), and digital assistants that schedule tasks that humans used to do. It may include also the 'computer therapist' ELIZA [10] and the mechanical arms in vehicle construction lines. They all have different requirements for HRI, as we shall discuss below.

3.1 Robots and Society in Sub-Saharan Africa

Where are robots used now, or could be used to alleviate problems? South Africa has about 28 industrial robots for each 10000 employees, which sits in between the other BRICS countries (3 in India and Russia, 10 in Brazil, and 68 in China) [6]. These robots are deployed mainly in manufacturing and mostly for cars, but may increase elsewhere, such as in food processing and sale, motivated by physical distancing due to COVID-19. Robots are relevant for the mining industry for, mainly, vertical tunnels, as drone reconnaissance, and rescue team assistance [15]. They may be robots or new 'cobots', which are collaborative robots that have a two-way interaction with the operator to complete the task [4,23].

Social robots, in contrast, are not intended to replace workers, but to provide services to make life easier. Healthcare is a popular use case topic for 'care robots' internationally, especially at the home to support independent living for the elderly [11]. Care for the elderly is of much lower relevance in Africa, due to customs in society and demographic makeup of the population with its median age of around 20 years. Rather, countries in Sub-Saharan Africa in particular have a much lower doctor/patient ratio than in the Global North and a substantial proportion of people live in remote locations with limited access to healthcare. This suggests care robots still could be useful, but they will have different uses and requirements. For instance, assistance with patient-doctor consultations, treatment instructions, medicine dispensing, after-care check-ups, and robot surgery, as ways to contribute to task-shifting healthcare work programmes. Since this involves sensitive data, ethics and solid implementation of privacy is crucial. Also, since there is less and slower internet access and less electricity, it will affect a robot's architecture and capabilities.

Another popular topic for societies, and that features prominent in the Sustainable Development Goals (SDGs), is education. There is a capability and capacity shortage in many African countries, although the data vary substantially per country. For instance, the learner/teacher ratio for primary school is 23.4 for the world, whereas within SADC countries, it varies between 23.7 (Botswana) and 58.7 (Malawi) with a SADC average of 38.8, and for secondary school, the world average is 17 versus a SADC average of 23.55, with again the lowest is Botswana (13.8) and the highest in Malawi (72.3)[iii]. Especially for the countries with a high learner/teacher ratio, such as Malawi, Angola, Mozambique, Tanzania, and Zambia, robots in education could alleviate the burden imposed on teachers. This could be for tasks such as homework assistance, automated marking and assignment feedback, as well as basic instruction in subjects for which there is a teacher shortage, such as home language, mathematics, and physics, as preferred over no assistance and uncertain future increase in number of teachers. Such uses fall within the identified categories of robot-supported teaching [28] and its soft-skills HRI capabilities requirements [9] that are yet to be operationalised. In this way, robots could assist in breaking the skills shortage vicious cycle and therewith uplift society, as compared to ethically problematic anti-social robot scenarios that threaten society's fabric with invasive surveillance and redundancies.

There are multiple other possibilities, such as the banking industry exploring digital assistants to broaden their reach into townships [7] and enhance bank transaction communication. Prolonged COVID-19 may push for robot scenarios that facilitate physical distancing. Autonomous weapons systems robots are not planned to be deployed [24].

3.2 Language Needs Analysis

As a first step in narrowing down the language needs, we categorise the type of robots and assess their primary modes of interaction, which is included in Table 1. Since speech technologies are becoming mainstream mainly for English only and typically still use text-to-speech, we will not consider it further[iv]. We zoom in on text, which is less hard computationally, less computationally intensive, and more results have been obtained.

Table 1. High-level categorisation of robot types and their predominant modes of interaction.

Robot category	**Mode of interaction**	**Comment**
Humanoid robot	Speech, possibly text	Device interaction is likely only speech, but if it is to behave human-like it also should be able to write and generate text
Non-humanoid robots, in society	Speech, else text	Primary mode of interaction with robots like a carebot or a barbot is speech; may have text (e.g., when a user has a speech impairment, is hard of hearing, or teaching a child to write)
Non-humanoid robots, in manufacturing	Multimodal	This may be natural language text, programming instructions, physical interaction (e.g., press a button, move joystick), and limited speech in the case of cobots
Software-based robots	Speech or text	These include the digital assistant devices, but also the text-based ones for assistance that pop up on websites

Considering the language interaction modes in Table 1, non-humanoid robots in society and software-based robots are more or less suitable for text-based interaction, which we shall focus on. It needs natural language *understanding* and *generation*. We narrow down the scope to text generation, known as Natural Language Generation (NLG), including Controlled Natural Languages (CNL). Many Sub-Saharan countries have a 'language of business' or 'language of government' following its colonial legacy—English, French, Portuguese, and Afrikaans—and one may expect that any robot deployment then would be in one of those languages. If a country has multiple official languages, the number of languages may increase; e.g., Zimbabwe has 16 official languages, South Africa 11, and Botswana and Malawi have a national language. The medium of communication for care and in education often occur in, and benefit from, the users' L1 language, which are predominantly in the NCB family of languages that are mostly agglutinating. Minimum requirements for text generation for HRI are then:

- Data-to-text generation; e.g., banking apps' messages about transactions and bank statements.
- Knowledge-to-text generation; e.g., Q/A about patient symptoms, treatment instructions, and language learning exercises and other subject matter engagement.
- Text generation for text-to-speech systems, which entails the requirement to use conversational style text (cf. written language).
- Text generation for end-user interaction, which requires sentence variability (cf. repetition) and a less formal style.
- Text generation in local languages, which are mainly of the NCB family, being at least the mother tongue (L1) language in healthcare settings and L1 and the language of teaching and learning in education.

While this may appear broad, it already excludes many common scenarios and technologies, such as CNLs for business rules management, chart-to-text, and creative writing NLG, and it entails the inability to use text generation techniques with sequence-to-sequence models because there is not enough online text in Niger-Congo languages. Also, most NLG systems assume text as the final output, so the type of outputs may have to be adjusted to what may sound good conversationally rather than what reads well.

4. Developments: Natural Language Generation for African languages

Examining NLG systems for NCB languages, let us briefly recap the traditional processing workflow for creating NLG systems. First, the *text planning* stage concerns the context what the NLG is used for, such as which HRI scenario with what sort of structured input. Second, *sentence planning,* to determine messages, words, and phrases. It may seem as if it can be done for English and simply be translated, but this is unlikely to be feasible due to sociolinguistic factors, as shown for isiZulu CNLs [17]. This is compounded in healthcare due to the practice of *ukuhlonipha*, i.e., using a different description of what ought not to be spoken of out loud, and there can be multiple variants of a language, which are not well-documented. Sentence planning may occur with people speaking one dialect and then be evaluated by participants speaking another, or have a mixed test group and obtain ambiguous results on key NLG evaluation criteria such as naturalness and grammaticality. Third, *linguistic realisation* has received most research attention to date. It concerns creating computational grammars gleaned from underspecified, outdated, and dispersed grammar documentation. The first NLG results were obtained for knowledge-to-text in isiZulu, including theory [18] and tooling [19], and subsequently isiXhosa [20] and Runyankore [2]. Well-resourced languages have off-the-shelf grammar engines, such as SimpleNLG for English [13] and for Portuguese [26] and French [29].

The next two sections zoom in on two use cases for HRI, following the pipeline: disease symptom description generation and computer-assisted language learning (CALL).

4.1 NLG for Healthcare in a Niger-Congo B language

The first use case is also a motivational scenario in related work [3,18]: communication in healthcare, of which language barrier issues, L1 benefits, and personalisation are well-known [5,16]. Patient disease and treatment summaries have to be generated, which should be able to be extracted and summarised automatically from electronic health records that are annotated with healthcare terminologies [5], notably SNOMED CT, and with OpenMRS and a suitable localisation in the desired language [27]. SNOMED CT is represented in the OWL language and thus requires knowledge-to-text NLG. A patient's record may state their symptoms of fever and a cough, which would be stored at the back-end as instances of the following axioms:

```
Subclassof(Patient ObjectSomeValuesFrom(hasSymptom Fever))
Subclassof(Patient ObjectSomeValuesFrom(hasSymptom Cough))
```

In the subsequent sentence planning stage, one could choose to apply sentence aggregation, since both axioms use the `hasSymptom` property. Aggregation is possible in English, but has not yet been addressed for any of the NCB languages. More importantly, a 'communication mode' has to be selected for the prospective application. Due to known language barrier issues that appear more frequently in Africa [16], we select *verification by the patient* to check that all relevant information of the consultation is recorded. Extant algorithms for any language can do this only in the rigid verbalisation mode for knowledge verification [18] rather than patient verification, such as with a template alike (in English for brevity)

Each [*class*] [*object property*] some [*class*] and some [*class*].

Extending the state of the art, a template or pattern for instance-level information, may be

[*individual*] [*object/data property*] some [*individual*] and some [*individual*]

to generate, respectively:

"Each Patient has as symptom some fever and some cough."

"Patient#12345 has as symptom some fever#567 and some cough#789."

Neither is suitable for the end-user scenario for application development that requires assistance in patient-doctor consultation, be it interaction, verification, or instructions. The AwezaMed tool [22] has text and speech capabilities, but it has canned text for a few illnesses only and only translates between English, Afrikaans, isiZulu and isiXhosa. What is needed, is flexibility and personalisation in the content, for any language. Thus, templates for less rigid sentences have to be developed, with different fixed-text and variable-text slots and in a conversational style; e.g., imagine the symptom verification interaction:

System: "You, registered as Patient#12345, have described to have as symptoms fever and cough. Is this correct?"

User: "Yes."

System: "Thank you for confirming this."

System reports to doctor: "The patient Patient#12345 confirms to have as symptoms fever and cough."

This has three consequences for system design. First, it means combining instance-level data with type-level knowledge to generate the sentence, which is unusual for ontology verbalisation. Second, it requires many more templates for the same representation of information so as to accommodate different contexts, and be interspersed with canned text sentences. Third, on the flip side, it eliminates the quantifiers from the sentences and therewith avoids a few cumbersome algorithms (since the surface realisation of the quantifiers are dependent on the noun class of the noun that it quantifies over). The latter consequence holds especially for all agglutinating NCB languages. For the former, a natural language-independent architecture may be devised.

Tools for generating treatment instructions induce a further requirement: while written instructions can be devoid of niceties, a (robot) doctor or task-shifting healthcare worker may speak friendly to a patient. Compare, e.g., "Take 3 pills daily" on a label with a robot dispenser telling the patient that "You must take 3 pills per day", "You should take 3 pills per day", or "I urge you, do take 3 pills per day". From a design viewpoint, this means creation and management of yet more template variants, and therewith more grammar rules to consider and more canned text to manage than just written instructions.

4.2 CALL Exercises for isiZulu and Related Agglutinating Languages

Typical international CALL questions and answers are about reading comprehension, vocabulary, and some grammar. This has been investigated especially for the widely spoken languages in the world and with varying levels of human intervention. Full automation for class assistance requires, as a minimum, the content input of exercise generation where the type of questions are limited but the vocabulary changes. Only few NLG and CNL-based attempts exist, for Latin [21], French [12], and isiZulu [14], where each one has a different scope: sentence writing assistance, question generation, and grammar, respectively. Overall, there are two NLG components to a CALL robot teacher: the exercises and the interaction.

Concerning the exercises, NLG can add variability and automated marking: instead of a teacher devising and marking, say, 25 singular-to-plural exercises, this is just one template with two sets of vocabulary (nouns and verbs) and the relevant grammar rule(s). The grammar rules for agglutinating languages and a noun class system—two interwoven features emblematic for NCB languages but not for languages outside of Africa—are more elaborate than regular expression based pluralisation rules. For instance, and inspired by the CALL proof-of-concept of [14], consider a template pair

Question: Convert the sentence to plural: `[noun_ncx_sg] [SC_ncx][verbstem]`

Answer expected and marked against: `[noun_ncy_pl] [SC_ncy][verbstem]`

and a set of nouns `{umuntu,ugogo,isilwane,inja}` and verb stems `{-dla,-lala,-phuza}`, then it can generate 12 exercises that can be automatically marked, including, e.g.:

Q: Convert the sentence to plural: ugogo udla
A: ogogo badla

Q: Convert the sentence to plural: inja idla
A: izinja zidla

These exercises generated are all unique, cf. repetitions in English ('the dog eats' and 'grandmother eats', resp.), thanks to the noun class system and concordial agreement. Also, the fixed text could be modified, by making "plural" a variable for mixing singular/plural exercises so that 24 unique exercises can be generated and automatically marked. The main tasks are matching sets of nouns and verbs, encode substantive grammar rules for more exercises, and difficulty levels. These CALL exercises are knowledge-to-text NLG.

The second component is the robot and its interaction. In the large classes of some 50-100 learners, a single classroom or telepresence type of robot is unrealistic, as are a robot per learner, but 4 or 5 robots to help manage homework in smaller groups might suffice. Since mobile phone penetration is high, but not network coverage, this could also be chatbot or digital assistant as an app on the smartphone. These scenarios require different conversational style interaction compared to the small classroom scenario in the Global North. First, it raises non-technical questions: how does HRI for groups differ from HRI for individuals, what amount of canned text can be used compared to templates or grammars, and to what extent is variation important versus repetition for learning tasks? The robot's requirements depend on what the expected tasks are; e.g., homework assistance or automatically generating and marking exercises, how feedback should be returned to the learner, whether it should include features of adaptive e-learning, and whether is it to offer content as well. For instance, regarding feedback, there can be CNL or NLG use, alike

Homework feedback template: You answered [x] out of [y] exercises correct. [Well done!{if x/y >= 0.8}/You're on the right track.{if 0.8 < x/y < 0.5}/Have a look at the grammar rules [first{if content count=0}/again{if content count >=1}] and let's try again.{if x/y =< 0.5}].

Exercise feedback template for an error: The plural of `[noun_ncx_sg]` is `[noun_ncy_pl]`, rather, which is in noun class `[ncy]` and its SC is `[SC_ncy]`.

The former focusses on data-to-text that is applicable anywhere, whereas the latter concerns knowledge-to-text tailored to the language. Thus, also in this use case, as for healthcare, an NLG requirement for the robots is to combine data-to-text and knowledge-to text, rather than one or the other that is more common with extant systems. In addition, it raises the requirement to resolve the number system: should the [x]'s be numbers or words? For instance, use, e.g., *ngu-10* (in isiZulu) or the appropriate completion of *-yishumi*, alike in

Uqede imisebenzi eziyishumi.

'You completed ten exercises.'

Numbers depend on what is counted over—e.g., '10 pills' is *amaphilisi ayishumi*—and thus are also governed by the noun class system. Currently, there are no algorithms for such numbers. Exploratory tests with basic numbers in Google Translate's English-isiZulu were promising, but its language models for isiZulu (and only 9 other African languages) are inaccessible and it requires internet, and therefore the service is impractical.

Finally, since all Sub-Saharan countries are multilingual, any CALL robot will have to be able to code switch between the language of teaching and learning, the learner's L1 (home/mother tongue) language, and the language that is being learned. Thus, the robot has to be multilingual and be able to code-switch, which puts forth the requirement to be able to detect the language spoken or written, for which only limited results are available [8].

5. Conclusions and Prospects

Social robots may be useful also in Sub-Saharan Africa, whose potential uses are at least partially distinct from social robots in the Global North. Key potential for deployment include assistance in large classes and as assistive remote primary healthcare. The Human-Robot Interaction is expected to be both in speech and text modes, rather than only speech, and will need to support African languages. A set of requirements were formulated for the natural language generation component of HRI by means of two small use cases. The lessons learned from them are the necessary key features of the mixed-mode of combining data- and knowledge-to-text in one system, a different sentence style for text-to-speech and consequently different grammar rules, verbalising numbers in a sentence, and code-switching between natural languages.

Although initial successes have been obtained in NLG and adjacent technologies for Niger-Congo B languages to work toward HRI, they are fragmented and are yet to be connected to complete at least one pipeline from input to dynamic, personalised, output in one language. A targeted proof-of-concept text-based HRI, such as a CALL roboteacher for grammar learning, may be available as early as within 2-3 years. Benefits to culturally appropriate HRI of relevant social robots in Africa may be manifold, since they are expected to alleviate long-term issues with skills shortages and therewith have the improvement of quality of life as main significance: better education of learners lifts their job prospects and the average level of work and care skills and a healthy breadwinner can provide better for their family. This while respecting their identities by means of interaction in one's own language. It also may inform NLG for HRI elsewhere, which is also in an exploratory phase as exhibited with its NLG4HRI2020 workshop and the Horizon 2020 project CARESSES. Follow-up projects could vary the scenarios and languages of interaction and focus on generic architectures to meet the elucidated requirements.

Acknowledgements This work was financially supported by the National Research Foundation (NRF) of South Africa (Grant Number 120852).

References

[1] Barnard, E., Davel, M. H. , van Heerden, C., et al. The NCHLT speech corpus of the South African languages. *Spoken Language Technologies for Under-Resourced Languages*, pp. 194-200, May 2014.
[2] Byamugisha, J. Ontology verbalization in agglutinating Bantu languages: a study of Runyankore and its generalizability. Department of Computer Science, University of Cape Town, South Africa. 2019.
[3] Byamugisha, J., Keet, C.M., DeRenzi, B. Evaluation of a Runyankore grammar engine for healthcare messages. *Proceedings of INLG'17*. ACL, pp105-113. 4-7 Sept 2017, Santiago de Compostela, Spain.

[4] Calitz, A.P., Poisat, P., Cullen, M. The future African workplace: The use of collaborative robots in manufacturing. *SA Journal of Human Resource Management*, 2017, 15(0), a901.
[5] Cawsey, J.A., Jones, B.R., Pearson, J. The evaluation of a personalized health information system for patients with cancer. *User Modeling and User-Adapted Interaction*, 2000, 10(1):47–72.
[6] Consultancy.co.za. South Africa's ratio of industrial robots to employees amongst the global top 40. Online: https://www.consultancy.co.za/news/654/south-africas-ratio-of-industrial-robots-to-employees-amongst-the-global-top-40; 4 June 2018. Last accessed: 2 Dec 2020.
[7] Davids, N. Touch down for SA's banking Siri. *UCT News*, 12 Sept 2019. Online: https://www.news.uct.ac.za/news/audio/-article/2019-09-12-touch-down-for-sas-banking-siri; Last accessed: 4 Dec 2020.
[8] Dube, M., Suleman, H. Language identification for South African Bantu languages Using Rank Order Statistics. *Proceedings of ICADL'19*, Springer LNCS 11853, 283-289.
[9] Edwards, B.I, Cheok, A.D. Why Not Robot Teachers: Artificial Intelligence for Addressing Teacher Shortage. *Applied Artificial Intelligence*, 2018, 32(4): 345-360.
[10] Weizenbaum, J. ELIZA--A Computer Program for the Study of Natural Language Communication Between Man and Machine. *Communications of the ACM*, 1966, 9(1): 36-42.
[11] Fischinger, D., Einramhof, P., Papoutsakis, K., et al., Hobbit, a care robot supporting independent living at home: First prototype and lessons learned. *Robotics and Autonomous Systems*, 2016, 75 (A), 60-78.
[12] Gardent, C., Perez-Beltrachini, L. Using FB-LTAG Derivation Trees to Generate Transformation-Based Grammar Exercises. *Proceedings of TAG+11*, 117-125, 2012.
[13] Gatt, A., Reiter, E. SimpleNLG: A realization engine for practical applications. *Proceedings of ENLG'09*, March 30-31, 2009, Athens, Greece. ACL, pp90–93.
[14] Gilbert, N., Keet, C.M. Automating question generation and marking of language learning exercises for isiZulu. *Proceedings of CNL'18*. IOS Press, 2018, FAIA vol. 304, 31-40.
[15] Green, J. Mine rescue robots requirements Outcomes from an industry workshop. *Proceedings of RobMech'13*, Durban, South Africa, 2013, pp. 111-116.
[16] Hussey, N. The language barrier: The overlooked challenge to equitable health care. *SAHR*, 189–195, 2012/13.
[17] Keet, C.M., Khumalo, L. Toward verbalizing logical theories in isiZulu. *Proceedings of CNL'14*. Springer LNAI, 8625, 78-89.
[18] Keet, C.M., Khumalo, L. Toward a knowledge-to-text controlled natural language of isiZulu. *Language Resources and Evaluation*, 2017, 51:131-157.
[19] Keet, C.M., Xakaza, M., Khumalo, L. Verbalising OWL ontologies in isiZulu with Python. *The Semantic Web: ESWC 2017 Satellite Events*. Springer LNCS, 10577, 59-64.
[20] Mahlaza, Z. Grammars for generating isiXhosa and isiZulu weather bulletin verbs. Department of Computer Science, University of Cape Town, South Africa. 2018.
[21] Lange, H., Ljunglöf, P. Putting Control into Language Learning. *Proceedings of CNL'18*. IOS Press FAIA vol 304, 61-70.
[22] Marais, L., et al. AwezaMed: A Multilingual, Multimodal Speech-To-Speech Translation Application for Maternal Health Care. *Proceedings of FUSION'20*, Rustenburg, SA, 2020, pp1-8.
[23] Masinga, P., Campbell, H. Trimble, J. A. A framework for human collaborative robots, operations in South African automotive industry. *Proceedings of IEEM'15*, Singapore, 2015, 1494-1497.
[24] Nxumalo. M. No 'killer robot' plans for SANDF. *IOL Daily News*, 3 Sept 2018. Online: https://www.iol.co.za/dailynews/news/no-killer-robot-plans-for-sandf-16827299; last accessed: 4 Dec 2020.
[25] Nthite, T., Tsoeu, M. End-to-End Text-To-Speech synthesis for under resourced South African languages. *Proceedings of SAUPEC/RobMech/PRAS'20*, Cape Town, South Africa, 2020, 6p. (abstract)
[26] Oliveira, R. de, Sripada, S. Adapting SimpleNLG for Brazilian Portuguese realisation. *Proceedings of INLG'14*, ACL, 93–94. Philadelphia, USA, 19-21 June 2014.
[27] Oluoch, T., de Keizer, N., Langat, P., *et al.* A structured approach to recording AIDS-defining illnesses in Kenya: A SNOMED CT based solution. *Journal of Biomedical Informatics*, 2015, 56: 387-394.
[28] Sharkey, A.J.C. Should we welcome robot teachers? *Ethics in Information Technology*, 2016,18:283-297.
[29] Vaudry, P-L., Lapalme, G. Adapting SimpleNLG for Bilingual English-French Realisation. *Proceedings of ENLG'13*, August 8-9, 2013, Sofia, Bulgaria. ACL, 2013, pp183-187.

---

[i] E.g., as promoted by the Deep Learning Indaba https://deeplearningindaba.com and by https://4irsa.org/
[ii] For instance, Robotic Innovations and Robotic Handling Systems for robots in manufacturing in South Africa, and RD-9 produces educational robots.
[iii] Data from the UNESCO Institute for Statistics (uis.unesco.org) as part of the SDG efforts; interactive graphs are available at https://data.worldbank.org/indicator/SE.PRM.ENRL.TC.ZS.
[iv] The interested reader is referred to extant efforts, such as [1, 25] for South African languages.