CSSR for

Culturally Sensitive Social Robotics for Africa

D5.4.1 Cultural Knowledge Ontology & Knowledge Base

Due date: **31/12/2024**
Submission Date: **30/12/2024**
Revision Date: **30/12/2024**

Start date of project: **01/07/2023** Duration: **36 months**

Lead organisation for this deliverable: **The University of the Witwatersrand**

Responsible Person: **D. Vernon** Revision: **1.0**

Project funded by the African Engineering and Technology Network (Afretec) Inclusive Digital Transformation Research Grant Programme		
Dissemination Level		
PU	Public	**PU**
PP	Restricted to other programme participants (including Afretec Administration)	
RE	Restricted to a group specified by the consortium (including Afretec Administration)	
CO	Confidential, only for members of the consortium (including Afretec Administration)	

Executive Summary

Deliverable D5.4.1 formalizes the Rwandan modes of social interaction documented in Deliverable D1.2: culturally sensitive behaviours, activities, actions, and motions, i.e., Rwandan cultural knowledge for polite and respectful interaction. It presents a cultural knowledge ontology and, based on this ontology, a simple representation of the cultural knowledge documented in D1.2, in the form of the cultural parameter values that can be used by the robot to emulate these polite and respectful behaviours, activities, actions, and motions.

In the work plan, this deliverable is assigned to the University of the Witswatersrand. However, the material in this version was developed and written by Carnegie Mellon University Africa. This was necessary because the ontology was needed to guide the preparation of the survey in Deliverable D1.2, and because a knowledge base was needed for integration in the system architecture described in Deliverable D3.1.

Contents

1 Introduction

Deliverable D5.4.1 formalizes the Rwandan modes of social interaction documented in Deliverable D1.2: culturally sensitive behaviours, activities, actions, and motions, i.e., Rwandan cultural knowledge for polite and respectful interaction. It presents a cultural knowledge ontology and, based on this ontology, a simple representation of the cultural knowledge documented in D1.2. This takes the form of the cultural key-value pairs that can be used by the robot to emulate polite and respectful behaviours, activities, actions, and motions.

We begin in Section 2 by addressing the representation of cultural knowledge, identifying the different categories of knowledge, and by introducing an ontology of cultural knowledge.

We then analyze the fifty-two consensus answers to the subset of the fifty-seven questions in the cultural knowledge survey documented in Deliverable D1.2, i.e., excluding questions 2-4, 2-5, 2-8, 3-28, and 3.30 which were discarded during the data cleaning phase. The goal of this analysis is to map each of the fifty-two questions and consensus answers to the ontology keys. The outcome of the analysis reveals that not all questions and consensus These are shown in Tables **??** and **??**; the questions that are excluded because they do not map to any of the ontology keys are marked with an asterisk. This resulted in a set of thirty-nine questions and associated answers, as shown in Table **??**.

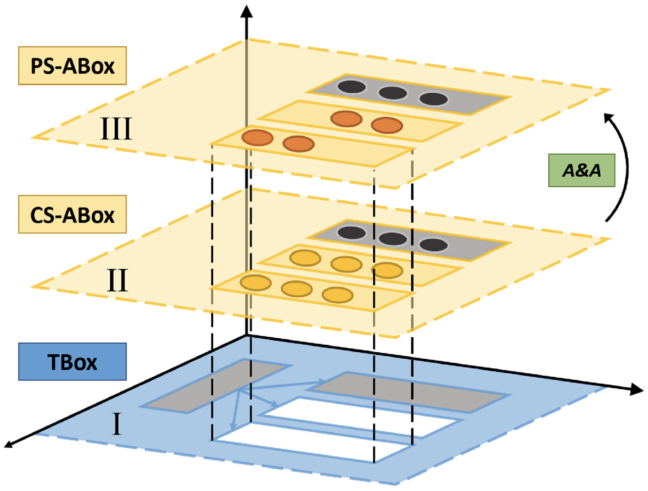

Figure 1: Knowledge representation architecture for a culturally competent robot. The bottom TBox layer (Layer I) defines the ontology for all knowledge, including domain-specific ontologies and upper ontologies that provide interoperability among domain-specific ontologies (grey boxes), and ontologies that model cultural-knowledge (white boxes). The middle CS-ABox layer (Layer II) is the culture-specific layer which includes instances of national-level cultural knowledge (yellow circles), as well as instances of knowledge from domain-specific ontologies and upper ontologies (grey circles). The top PS-ABox layer (Layer III) is the person-specific layer which includes instances of knowledge about the user (orange circles), as well as instances of knowledge from domain-specific ontologies and upper ontologies (grey circles). (From Bruno et al. 2019 [1].)

2 Representation of Cultural Knowledge

In the following, we summarize the knowledge representation architecture and knowledge classification suggested by Bruno et al. [1] and explain how adopting elements of this classification facilitates the creation of a knowledge representation that can be used in the CSSR4Africa system.

We then explain how the answers to each question in the survey can be mapped to the ontology. We then define a simple representation of the knowledge using key-value pairs, with keys derived from the ontology.

2.1 The Different Categories of Knowledge

Bruno et al. [1] propose a knowledge representation architecture for a culturally competent robot; see Fig. 1. This architecture has three layers, each capturing a different element of the knowledge specification. The bottom layer is a *terminological* box (TBox). This is where the ontology proper is specified. The middle and top layers are *assertional* boxes (ABox). This is where the culture-specific and person-specific knowledge (defined by the ontology) is stored.

In more detail, the three elements of the knowledge representation architecture are as follows.

A culture-generic knowledge ontology is captured in the bottom TBox layer (Layer I). This layer defines the ontology for all knowledge, including domain-specific ontologies and upper ontologies that provide interoperability among domain-specific ontologies (grey boxes), and ontologies that model cultural-knowledge (white boxes).

Culture-specific knowledge is captured in the middle CS-ABox layer (Layer II). Specifically, this layer includes instances of national-level cultural knowledge (yellow circles), as well as instances of knowledge from domain-specific ontologies and upper ontologies.

Person-specific knowledge is captured in the top PS-ABox layer (Layer III), including instances of knowledge about the user (orange circles), as well as instances of knowledge from domain-specific ontologies and upper ontologies

The culture-generic knowledge ontology captures eight types of knowledge, grouped in three categories, as follows.

1. Context knowledge.
   (a) Knowledge about the assisted person.
   (b) Knowledge about the environment.
2. Robot knowledge.
   (a) Knowledge about the actions that the robot can perform.
   (b) Knowledge about the parameters of these actions.
   (c) Knowledge about how actions can be combined into higher-level behaviours.
3. Core values knowledge.
   (a) Knowledge about the goals of the robot mission.
   (b) Knowledge about social norms; these can be considered additional culturally-grounded goals, i.e., constraints on goals, planning operators, action, and cultural parameters.
   (c) Knowledge about conversational subject matter.

Here, we are concerned with 2 (a) knowledge about the actions that the robot can perform, 2 (b) knowledge about the parameters of these actions, and 3 (b) knowledge about social norms. The values we use for the action and cultural parameters determine the culturally sensitive nature of the robot's actions. To quote Bruno et al. [1]:

> "Knowledge pertaining to the robot's sensorimotor and communication capabilities is required by the robot to know what it can do and how the user might prefer it to be done. This knowledge again includes static, *a priori* information (e.g., describing the set of commands allowing the robot to perform the Namaste greeting, the associated parameters and their preferable values) and dynamic information (e.g., describing the robot's current posture and values of related parameters)."

These values are then used by the various ROS nodes in the CSSR4Africa system when invoking actions through ROS service requests. The values themselves are derived from the consensus answers to the survey questions.

The actions that the robot can perform — 2 (a) — depend on the functionality of the system architecture, as described in Deliverable D3.1: animate behaviour, deictic, iconic, and symbolic gesture, overt attention, locomotion and navigation. As such, we do not encode this knowledge explicitly in the CSSR4Africa knowledge base. Neither do we encode knowledge about how actions can be combined into higher-level behaviours — 2 (c) — explicitly in the CSSR4Africa knowledge base,

although some of the knowledge that is revealed and made explicit by the consensus answers to the survey questions does address activities and behaviours. Thus, the cultural knowledge that is required for the CSSR4Africa project has two forms:

1. A compendium of culturally sensitive behaviours, activities, actions, and motions.
2. An knowledge ontology to categorize the behaviours, activities, actions, and motions that the Pepper robot can perform.
3. A mapping from the compendium of culturally sensitive behaviours, activities, actions, and motions to the ontology.
4. The action and cultural parameter values — 2 (b) and 3 (b) — that are derived from a subset of the consensus answers.

Item 1 comprises the consensus answers to the fifty-seven questions in the Rwandan cultural knowledge survey documented in Deliverable D1.2. The remaining items are documented in this deliverable. Note that the mapping from the compendium of culturally sensitive behaviours, activities, actions, and motions to the ontology is partial because there are some behaviours, activities, actions, and motions that the Pepper robot is incapable of performing.

2.2 A Knowledge Ontology for the Pepper Robot

As noted in Section 1, Deliverable D1.2 compiles the cultural knowledge required for culturally sensitive human robot interaction between robots and Rwandan people. To be effective, this knowledge must be organized in some manner. This organization is effectively the knowledge ontology that is the result of in Task 5.4.1, and documented in this deliverable in Figure 2.

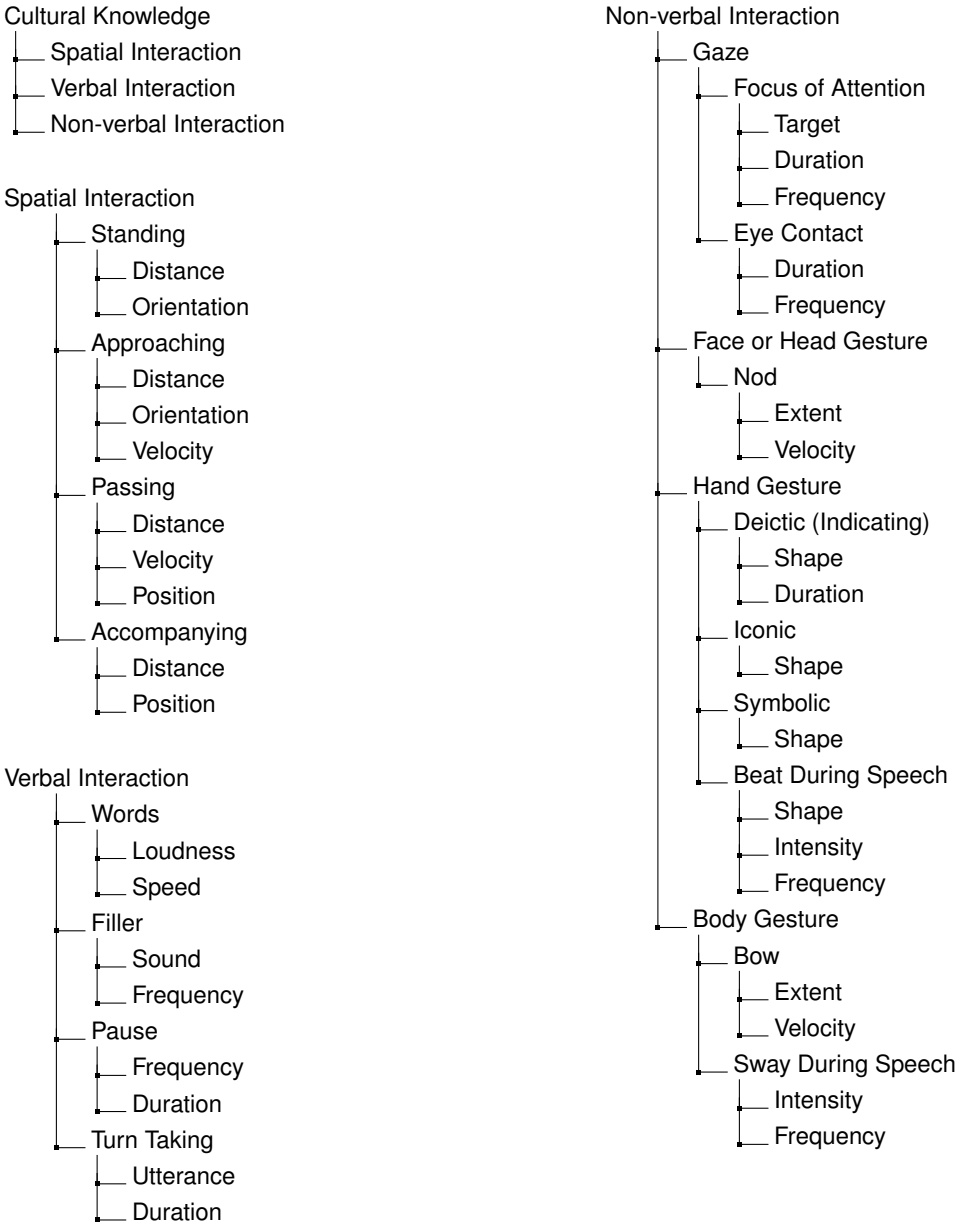

Figure 2: Ontology of Cultural Knowledge. Note that the ontology is restricted to the actions that the Pepper robot can perform.

3 Mapping Rwandan Cultural Knowledge to the Knowledge Ontology Keys

By itself, an ontology is insufficient, since the knowledge base that is to be used by the robot must be populated by the knowledge that is derived from the survey and documented in Deliverable D1.2. Therefore, we also need to map the knowledge in the answers to each question in the survey to the ontology.

Tables 1 and 2, reproduced from Deliverable D1.2, presents the consensus cultural knowledge produced by the survey described in D1.2. These capture the behaviors, activities, actions, and motions that are considered polite and respectful when interacting with people in Rwanda. Here, we have identified those questions that do not map directly to the knowledge ontology keys by appending an asterisk to the question number. In the next section, we provide an explicit mapping between the questions without asterisks (and the associated consensus cultural knowledge) and the knowledge ontology.

Question	Consensus Cultural Knowledge
2-1	To show respect, one should lower gaze when greeting someone older.
2-2	One should suspend work or movements and pay attention when addressed.
2-3	One should keep intermittent eye contact; lack of eye contact depicts disrespect as it shows divided attention during the interaction
2-6	One should use an open palm of the hand to point to people and objects.
2-7	One should not point an upward facing palm of the hand at someone.
2-9	To show respect, one should bow slightly when greeting someone older.
2-10	To show respect, one should raise both hands when greeting.
2-11 *	One should not wave at someone from a distance; one should move towards them to greet them
2-12 *	One should not use the left hand to hand something to someone
2-13 *	To show respect, one should hand over and accept gifts with two hands and do so from the front, facing the recipient
2-14 *	To show respect, one should shake hands with the right hand and use the left arm to support the right forearm when doing so.
2-15 *	An appreciation of rhythmic sound and movement is valued.
2-16	To show respect, one should bow slightly and lower gaze when greeting someone older
2-17	The younger interaction partner should bow when greeting an older person or when rendering a service
2-18 *	All interactions should begin with a courteous greeting.
2-19	The younger interaction partner should enable a greeting to be initiated by an older person.
2-20 *	It is respectful to use local languages and they should be used for verbal interaction when possible.
2-21 *	One should use formal titles when addressing someone.
2-22 *	One should engage in a preamble before getting to the point, as being too forward may be regarded as disrespectful.
2-23	One should not interrupt or talk over someone when they are speaking.
2-24	One should not talk loudly to an older person
2-25 *	Behaviours should focus on fostering social connections and relationships; they should not be purely functional.
2-26	One should not walk between two or more people who are conversing because it is considered rude to do so.
2-27	One should not walk far ahead of an older person, unless leading the person (in which case, one should walk slightly to the side).

Table 1: Consensus answers to the subset of the twenty-seven questions in Part 2 of the cultural knowledge survey. Answers to questions 2-4, 2-5, and 2-8 are not listed as no consensus could be identified. Questions marked with an asterisk can not be mapped to the knowledge ontology.

Question	Consensus Cultural Knowledge
3-1	One should maintain a distance of one meter or less when passing someone.
3-2	One should say 'Hello' or 'Muraho' when acknowledging someone while passing them.
3-3	One should pass behind a group of two or more people.
3-4	One should position themselves beside someone older when showing them the way.
3-5	One should position themselves beside someone of the same age when showing them the way.
3-6	One should position themselves beside someone younger when showing them the way.
3-7–3-9 *	The preferred way to address someone, whether they are older, younger, or the same age, and whom you haven't met before, is by saying 'Muraho' or 'Hello'.
3-10	When asked a question, respondents should pause for a few seconds before answering.
3-11	In turn-based conversations, participants can raise their right hand to signal their desire to speak.
3-12	When explaining something to someone, you should direct your gaze equally between the person and the object.
3-13	When explaining something to someone, you should make eye contact often.
3-14	You should make eye contact more often when explaining something to someone older than you.
3-15	You should make eye contact more often when explaining something to someone younger than you.
3-16	When someone is explaining something to you, you should direct your gaze equally between the person and the object.
3-17	When someone is explaining something to you, you should make eye contact often.
3-18	If someone is explaining something to you and they are older than you, you should make eye contact more often.
3-19	If someone is explaining something to you and they are younger than you, you should make eye contact more often.
3-20	To draw someone's attention to something, use a head-nodding gesture while looking at the object.
3-21	To express gratitude, common gestures include nodding, smiling, and bowing the head, using hand gestures like a thumbs up or clasped hands, and slight bowing of the body.
3-22	To express agreement, common gestures include nodding the head and giving a thumbs up with the right hand.
3-23	To show respect, common gestures include a slight bow of the head, a greeting or handshake using the right hand supported by the left, and bowing, which is the most frequent body gesture.
3-24	To express friendliness, people commonly use facial gestures like smiling, hand gestures such as a handshake using both hands or the right hand, and body gestures like hugging.
3-25	When expressing confusion, individuals typically use facial gestures like wrinkling or frowning the brow or tilting the head, hand gestures such as raising both hands or the right hand, and body movements that vary according to the situation.
3-26	When expressing comprehension, individuals typically use head gestures, such as nodding, hand gestures like a right-hand thumbs-up, and body gestures that vary by situation.
3-27	When expressing interest, nodding and smiling are the most common gestures, while hand gestures such as giving a thumbs up with the right hand and body gestures like facing someone are used less frequently.
3-29	One should use body and hand gestures while speaking to someone, which depends on the situation. The most recommended gestures are slight body movement and slightly moving both hands.

Table 2: Consensus answers to the subset of the thirty questions in Part 3 of the cultural knowledge survey. Answers to questions 3-28 and 3-30 are not listed as no consensus could be identified. Questions marked with an asterisk can not be mapped to the knowledge ontology.

4 Knowledge Representation

A way of representing the knowledge presented in the previous section is also required. We propose a simple knowledge representation based on the knowledge categories suggested by Bruno et al. [1]. This requires the cultural knowledge ontology in Figure 2 to align with the parameters of the robot actions, as suggested by Bruno et al. [1].

While Bruno et al. [1] use the OWL-2 language to define their ontology, we adopt a simpler approach here that represents the ontology as a tree of concepts, as shown in Figure 2. Note that the ontology is restricted to the actions that the Pepper robot can perform. It explicitly omits forms of non-verbal communication that are important in human-robot interaction, e.g., facial expressions, such as eyebrow and mouth gestures.

This provides us with a straightforward way to specify the parameter values for each element in the ontology: we can represent the cultural knowledge with a simple list of key-value pairs, where a key is constructed from the name of a leaf nodes in the ontology tree and the name of its parent. The values can be either quantitative numeric values or qualitative symbolic values, which can then be interpreted by the ROS node that uses the key-value pair to produce culturally sensistive behaviour. If the survey answers require more than one value for a key, this can be accommodated by appending a suffix to the key, e.g., `SymbolicShapeWelcome` instead of `SymbolicShape`, and updating the ontology accordingly.

Table 3 lists the keys derived from the ontology tree, and identifies the questions in Parts 2 and 3 of the survey that reveal the associated cultural knowledge, and listed in Tables 1 and 2.

Tables 4–6 lists the keys derived from the subtrees of the ontology tree — Spatial Interaction, Verbal Interaction, and Non-verbal Interaction — and, in the next version of this deliverable, it will identify the numeric and symbolic values that can be associated with these keys based on the answers to the questions in the survey.

Key	Questions
Spatial Interaction	
StandingDistance	
StandingOrientation	
ApproachingDistance	
ApproachingOrientation	
ApproachingVelocity	
PassingDistance	3-1
PassingVelocity	
PassingPosition	2-26, 3-3
AccompanyingDistance	2-27, 3-4, 3-5, 3-6
AccompanyingPosition	
Verbal Interaction	
WordLoudness	2-24
WordAddressMethod	2-24
WordSpeed	
FillerSound	
FillerFrequency	
PauseFrequency	
PauseDuration	3-10
TurnTakingUtterance	2-19, 2-23 , 3-11
TurnTakingDuration	
Non-Verbal Interraction	
Gaze	
FocusofAttentionTarget	2-1, 2-2, 2-16, 3-12, 3-16
FocusofAttentionDuration	
FocusofAttentionFrequency	
EyeContactDuration	2-3, 2-4, 2-5, 3-2
EyeContactFrequency	3-13, 3-14, 3-15, 3-17, 3-18, 3-19
Face or Head Gesture	
NodExtent	3-2, 3-20, 3-21, 3-22, 3-23, 3-24, 3-25, 3-26, 3-27
BowExtent	3-24
FacialGesture	3-24
NodVelocity	3-2, 3-20, 3-21, 3-22, 3-23, 3-24, 3-25, 3-26, 3-27
Hand Gesture	
DeicticShape	2-6, 2-7, 2-8
DeicticDuration	
IconicShape	3-29, 3-30
SymbolicShape	2-10, 3-2, 3-21, 3-22, 3-23, 3-24, 3-25, 3-26, 3-27
BeatShape	
BeatIntensity	
BeatFrequency	
Body Gesture	
BowExtent	2-9, 2-16, 2-17, 3-21, 3-22, 3-23, 3-24, 3-25, 3-26, 3-27
BowVelocity	
SwayIntensity	3-21, 3-22, 3-23, 3-24, 3-25, 3-26, 3-27
SwayFrequency	
IconicShapeSpeaking	3-29

Table 3: Keys for specifying culturally sensitive actions and the questions in Parts 2 and 3 of the survey that reveal the associated cultural knowledge.

Key	Values
Spatial Interaction	
StandingDistance	
StandingOrientation	
ApproachingDistance	
ApproachingOrientation	
ApproachingVelocity	
PassingDistance	1m or less
PassingVelocity	
PassingPositionAvoid	Walk between two or more people who are conversing
PassingPositionPreferred	Pass behind a group of two or more people
AccompanyingDistanceAvoid	Not walk far ahead of an older person
AccompanyingDistancePereffered	Position beside someone when showing them the way
AccompanyingPosition	

Table 4: Spatial interaction: key-value pairs for specifying culturally sensitive actions.

Key	Values
Verbal Interaction	
WordLoudness	Do not talk loudly to an older person
WordAddressMethod	‘Muraho‘ or ‘Hello‘
WordSpeed	
FillerSound	
FillerFrequency	
PauseFrequency	
PauseDuration	A few seconds
TurnTakingUtteranceSignal	Raise the right hand
TurnTakingUtteranceAvoid	Do not talk over someone when they are speaking
TurnTakingUtteranceInitiates	The younger partner should let the older person initiate the greeting
TurnTakingDuration	

Table 5: Verbal interaction: key-value pairs for specifying culturally sensitive actions.

Key	Values
Non-Verbal Interaction	
Gaze	
FocusofAttentionTargetGreetingOlder	Lower your gaze and bow slightly
FocusofAttentionTargetAddressed	When addressed, stop what you are doing and give your full attention
FocusofAttentionTargetExplanation	Direct your gaze equally between the person and the object
FocusofAttentionDuration	
FocusofAttentionFrequency	
EyeContactDurationInteraction	Keep intermittent eye contact
EyeContactFrequencyExplain	Make eye contact often
EyeContactFrequencyExplainOlder	Make eye contact often
EyeContactFrequencyExplainYounger	Make eye contact often
EyeContactFrequencyListen	Make eye contact often
EyeContactFrequencyListenOlder	Make eye contact often
EyeContactFrequencyListenYounger	Make eye contact often
Face or Head Gesture	
NodExtentAttention	Use a head-nodding gesture while looking at the object
BowExtentGratitude	Slight bow of the head
NodExtentAgreement	Slightly nodding of the head
BowExtentRespect	Slight bow of the head
FacialGestureFriendliness	Smiling
FacialGestureConfusion	Wrinkling or frowning the brow or tilting the head
NodExtentComprehension	Nodding the head
NodExtentListening	Nodding the head
NodVelocity	
Hand Gesture	
DeicticShapePoint	Use an open palm to point at people and objects
DeicticShapePointAvoid	Do not point an upward-facing palm at someone
DeicticDuration	
IconicShapeSpeaking	Slightly moving both hands
SymbolicShapeRespect	Raise both hands when greeting.
SymbolicShapeRespectHandShake	Handshake with the right hand, supported by the left, while bowing
SymbolicShapeGratitude	Thumbs-up or clasped hands
SymbolicShapeAgreement	Give a thumbs-up with the right hand
SymbolicShapeFriendliness	Handshake with both hands or just the right hand
SymbolicShapeConfusion	Raise hand
SymbolicShapeComprehension	Right-hand thumbs-up
SymbolicShapeAvoid	Pointing a finger
SymbolicShapeAvoidGreeting	Greeting with the left hand
BeatShape	
BeatIntensity	
BeatFrequency	
Body Gesture	
BowExtentGreeting	Bow slightly when greeting someone older.
BowExtentGratitude	Bow slightly
BowExtentRespect	Bow slightly
BowVelocity	
SwayIntensity	
SwayFrequency	
IconicShapeSpeaking	Slight body movement

Table 6: Non-verbal interaction: key-value pairs for specifying culturally sensitive actions.

References

[1] B. Bruno, C. T. Recchiuto, I. Papadopoulos, A. Saffiotti, C. Koulouglioti, R. Menicatti, F. Mastrogiovanni, R. Zaccaria, and A. Sgorbissa. Knowledge representation for culturally competent personal robots: requirements, design principles, implementation, and assessment. *International Journal of Social Robotics*, 11(3):515–538, 2019.

Principal Contributors

The main authors of this deliverable are as follows (in alphabetical order).

Eyerusalem Birhan, Carnegie Mellon University Africa.
Muhirwa Richard, Carnegie Mellon University Africa.
David Vernon, Carnegie Mellon University Africa.

Document History

Version 1.0

First version created by moving and reorganizing Section 2 Representation of Cultural Knowledge and Section 3.2 Action and Cultural Parameter Values fom Deliverable D1.2.
David Vernon.
30 December 2024.