

Conversational Homes

Nick O’Leary, Dave Braines

Emerging Technology,

IBM United Kingdom Ltd,

Hursley Park, Winchester, UK

Email: nick_oleary|dave_braines@uk.ibm.com

Alun Preece, Will Webberley

School of Computer Science and Informatics,

Cardiff University, Cardiff, UK

Email: PreeceAD|WebberleyWM@cardiff.ac.uk

Abstract—As devices proliferate, the ability for us to interact with them in an intuitive and meaningful way becomes increasingly challenging. In this paper we take the typical home as an experimental environment to investigate the challenges and potential solutions arising from ever-increasing device proliferation and complexity. We show a potential solution based on conversational interactions between “things” in the environment where those things can be either machine devices or human users. Our key innovation is the use of a Controlled Natural Language (CNL) technology as the underpinning information representation language for both machine and human agents, enabling humans and machines to trivially “read” the information being exchanged. The core CNL is augmented with a conversational protocol enabling different speech acts to be exchanged within the system. This conversational layer enables key contextual information to be conveyed, as well as providing a mechanism for translation from the core CNL to other forms, such as device specific API requests, or more easily consumable human representations. Our goal is to show that a single, uniform language can support machine-machine, machine-human, human-machine and human-human interaction in a dynamic environment that is able to rapidly evolve to accommodate new devices and capabilities as they are encountered.

Keywords—IoT; Controlled Natural Language; Conversational Interaction.

I. INTRODUCTION

From an individual agent’s perspective, the Internet of Things (IoT) can be seen as an increasingly large and diverse world of other agents to communicate with. Humans are agents too in this world, so we can observe four kinds of communication: (i) human-machine, (ii) machine-human, (iii) machine-machine, and (iv) human-human. There is a tendency to consider human-oriented (i, iv) and machine-oriented (ii, iii) interactions as naturally requiring different kinds of communication language; humans prefer natural languages, while machines operate most readily on formal languages. In this paper, however, we consider what the IoT world might look like where humans and machines largely use a common, uniform language to communicate. Our design goal is to support communication activities such as: the discovery of other agents and their capabilities, querying other agents and receiving understandable information from them, and obtaining rationale for an agent’s actions. The proposed approach must be able to cope with rapid evolution of an IoT environment that needs to accommodate new devices, capabilities, and agent types. In Section II, we consider why human users might find such an environment more appealing when machines communicate using an accessible and human-friendly language, than when machines use a traditional machine-to-machine formalism. Section III substantiates our proposed

approach using a series of vignettes, while section IV provides some initial experimental evidence that human-machine and machine-machine interactions can be facilitated via a CNL communication mechanism. Section V concludes the paper.

II. BACKGROUND AND RELATED WORK

A key part of our approach is to consider the way in which humans “want” to interact with machines in the world. To help us gain insights into these latent human requirements we look towards existing trends and events occurring in the world and use these as inspiration to help us form our hypotheses about what a conversational environment for human-machine agents might entail. For example, in this work we consider recent interest in conversational technologies such as chatbots [1], conversational computing [2], and conversational agents [3]. The remainder of this section covers this human-motivated perspective and develops ideas first presented in [4].

A. Social Things

The advent of Twitter as a means of social communication has enabled a large number of otherwise inanimate objects to have an easily-accessible online presence. For example, Andy Stanford-Clark created an account for the Red Funnel ferries that service the Isle of Wight in the UK. The account [5] relays real-time information about the ferry arrivals and departures allowing a subscriber of the account to see if they are running on time.




Figure 1: Redjet tweet example

Another similar example is an unofficial account for London’s Tower Bridge [6]. Its creator, Tom Armitage, created a system that took the published scheduled of bridge opening and closing times and produced a Twitter account that relayed that information.



Figure 2: Tower Bridge tweet example

A key difference between the ferries and the bridge accounts is that the ferries are just relaying information, a

timestamp and a position, whereas the bridge is speaking to us in the first-person. This small difference immediately begins to bring a more human nature to the account. But, they are ultimately simple accounts that relay their state to whomever is following them, providing an easily consumable feed of information on an existing platform.

This sort of thing seems to have caught on particularly with the various space agencies. We no longer appear able to send a robot to Mars, or land a probe on a comet without an accompanying Twitter account bringing character to the events. The Mars Curiosity Rover has had an account [7] since July 2008 and regularly shares images it has captured. There's always a sense of excitement when these inanimate objects start to have a conversation with one another. The conversations between the European Space Agency Philae lander [8] and its Rosetta orbiter [9], as the former began to lose power and had to shutdown, generated a large emotional response on social media. The lander, which was launched into space years before social media existed, chose to use its last few milliamps of power to send a final goodbye.

The reality, of course, is that the devices did not create these tweets. Communication with them remains the preserve of highly specialized engineers, and their personalities are a creation of their public relations agencies on this planet. There are however, examples of machine participation on social media provided by social bots [10]. On occasion, these entities can masquerade as human agents and alter the dynamics of social sense-making and social influence.

B. Seamlessness vs Seamfulness

The IoT makes possible a future where our homes and workplaces are full of connected devices, sharing their data, making decisions, collaborating to make our lives better [11]. Whilst there are people who celebrate this invisible ubiquity and utility of computing, the reality is going to be much more complicated.

Mark Weiser, Chief Scientist at Xerox PARC, coined the term “ubiquitous computing” in 1988 as recognition of the changing nature of our interaction with computers [12]. Rather than the overt interaction of a user sitting in front of a computer, ubiquitous computing envisages technology receding into the background of our lives.

Discussion of ubiquitous computing often celebrates the idea of seamless experiences between the various devices occupying our lives. Mark Weiser advocated for the opposite; that seamlessness was undesirable and a self-defeating attribute of such a system. He preferred a vision of “Seamfulness, with beautiful seams” [13].

The desire to present a single view of a system, with no joins, is an unrealistic aspiration in the face of the cold realities of Wi-Fi connectivity, battery life, system reliability and the status of cloud services. Presenting a user with a completely monolithic system gives them no opportunity to connect with, and begin to understand, the constituent parts. That is not to say all users need this information all of the time, but there is clearly utility to some users some of the time: when you come home from work and the house is cold, what went wrong? Did the thermostat in the living room break and decide it was the right temperature already? Did the message from the working thermostat fail to get to the boiler? Is the boiler broken? Did

you forget to cancel the entry in your calendar saying you'd be late home that day? Without some appreciation of the moving parts in a system, a user cannot feel any ownership or empowerment when something goes wrong with it. Or worse yet, how can they avoid feeling anything other than intimidated by this monolithic system that responds in a manner akin to, “I'm Sorry Dave, I'm afraid I can't do that”.

This is the justification for beautiful seams: they help you understand the edges of a device's sphere of interaction, but should not be so big as to trip you up. For example, such issues exist with the various IP connected light bulbs that are available today. When a user needs to remember which application to launch on their phone depending on which room they are walking into and which manufacturer's bulbs happen to be in there, the seams have gotten too big and too visible.

Designer Tom Coates has written on these topics [14]. He suggests the idea of having a chat-room for the home:

“Much like a conference might have a chat-room so might a home. And it might be a space that you could duck into as you pleased to see what was going on. By turning the responses into human language you could make the actions of the objects less inscrutable and difficult to understand...”

This relates back to the world of Twitter accounts for Things, but with a key evolution. Rather than one-sided conversations presenting raw data in a more consumable form, or Wizard-of-Oz style man-behind-the-curtain accounts, a chat-room is a space where the conversation can flow both ways; both between the owner and their devices, and also between the devices themselves.

C. Getting Things Communicating

For devices to be able to communicate they need to share a common language. Simply being able to send a piece of data across the network is not sufficient. As with spoken language, the context of an interaction is important too.

This model of interaction applies to both the data a device produces, as well as the commands it can consume. There are a number of technologies available for producing such a shared model. For example: HyperCat [15], a consortium of companies funded by the UK Government's innovation agency in 2014. It provides a central catalog of resources that are described using RDF-like triple statements. Each resource is identified by a URI allowing for ease of reference. URIs are a key component in building the World Wide Web and are well understood, but they are a technology used primarily by computers. They do not provide a human-accessible view of the model.

Furthermore, to enable a dynamic conversation, any such model needs to be adaptable to the devices that are participating, especially when one of those participants is a human being.

D. Talking to Computers

The most natural form of communication for most humans is that of their own spoken language, not some JSON or XML encoded format that was created with the devices as the primary recipient. Technical specialists can be trained to understand and use technical machine languages, but this overhead is not acceptable to more casual everyday users who may wish to interact with the devices in their home. In addition

to this, we are living in an age where talking to computers is becoming less the preserve of science fiction: Apple's Siri, OK Google, Microsoft Cortana all exist as ways to interact with the devices in your pocket. Amazon Echo exists as a device for the home that allows basic interaction through voice commands. This means that there is now a plausible expectation that an everyday person could interact with complex devices in their home in a natural conversational manner.

Natural Language Processing (NLP) is one of the key challenges in Computer Science [16]. In terms of speech understanding, correctly identifying the words being spoken is relatively a well-solved problem, but understanding what those words mean, what intent they try to convey, is still a hard thing to do.

To answer the question "Which bat is your favorite?" without any context is hard to do. Are we talking to a sportsperson with their proud collection of cricket bats? Is it the zookeeper with their colony of winged mammals? Or perhaps a comic book fan is being asked to choose between incarnations of their favorite super hero.

Context is also vital when you want to hold a conversation. Natural language (NL) is riddled with ambiguity. Our brains are constantly filling in gaps, making theories and assumptions over what the other person is saying. For humans and machines to communicate effectively in any such conversational home setting, it is important that contextual information can be communicated in a simple, but effective, manner. This must be achieved in a manner that is accessible to both the human and machine agents in this environment.

III. CONTROLLED NATURAL LANGUAGE

To avoid a lot of the hard challenges of NLP, a CNL can be used. A CNL is a subset of a NL that uses a restricted set of grammar rules and a restricted vocabulary [17]. It is constructed to be readable by a native speaker and represents information in a structured and unambiguous form. This also enables it to be read and properly interpreted by a machine agent via a trivial parsing mechanism without any need for complex processing or resolution of ambiguity. CNLs range in strength from weaker examples such as simple style guides, to the strongest forms that are full formal languages with well-defined semantics. In our work, to identify a unifying language for both human and machine communication, we are focused on languages at the strong end of the scale, but we additionally wish to retain the requirement for maximal human consumability.

Ambiguity is a key issue for machine agents: whilst human readers can tolerate a degree of uncertainty and are often able to resolve ambiguity for themselves, it can be very difficult for a computer to do the same. CNLs typically specify that words be unambiguous and often specify which meaning is allowed for all or a subset of the vocabulary. Another source of ambiguity is the phrase or sentence structure. A simple example is concerned with noun clusters. In English, one noun is commonly used to modify another noun. A noun phrase with several nouns is usually ambiguous as to how the nouns should be grouped. To avoid potential ambiguity, many CNLs do not allow the use of more than three nouns in a noun phrase.

There are two different philosophies in designing a CNL. As mentioned previously a weaker CNL can be treated as a

simplified form of NL with a stronger CNL as an English version of a formal language [18]. In the case of a simplified form of NL, it can allow certain degrees of ambiguity in order to increase human accessibility. It relies on standard NLP techniques, lexical-semantic resources and a domain model to optimize its interpretation.

The alternative is to treat a CNL as an entirely deterministic language, where each word has a single meaning and no ambiguity can exist. Whilst computationally very efficient, it can be hard for a human user unfamiliar with the particular lexicon and grammar to write it effectively. This is because it competes with the user's own intuition of the language. The closer a CNL is to corresponding NL, the more natural and easy it is to use by humans, but it becomes less predictable and its computational complexity increases. The converse is also true. The more deterministic the CNL is, the more predictable it is, but the more difficult it is for humans to use.

In summary, in the operational setting described in this paper a CNL is designed to support both human usage and machine processing. It provides:

- 1) A user-friendly language in a form of English, instead of, for example, a standard formal query language (such as SPARQL or SQL). Enabling the user to construct queries to information systems in an intuitive way.
- 2) A precise language that enables clear, unambiguous representation of extracted information to serve as a semantic representation of the free text data that is amenable to creating rule-based inferences.
- 3) A common form of expression used to build, extend and refine domain models by adding or modifying entities, relations, or event types, and specifying mapping relations between data models and terminology or language variants.
- 4) An intuitive means of configuring system processing, such as specifying entity types, rules, and lexical patterns.

A good balance between the naturalness and predictability of the CNL is fundamentally important, especially to the human users as the strength and formality of the language increases.

A. An Introduction to ITA Controlled English

In previous research, we have proposed a specific CNL that is a variant of "Controlled English" known as ITA Controlled English, or just "CE" in shorthand [19]. This has been researched and developed under the International Technology Alliance (ITA) in Network and Information Science [20]. CE is consistent with First Order Predicate Logic and provides an unambiguous representation of information for machine processing. It aspires to provide a human-friendly representation format that is directly targeted at non-technical domain-specialist users (such as military planners, intelligence analysts or business managers) to enable a richer set of reasoning capabilities [21] [22].

We assert that CE can be used as a standard language for representation of many aspects of the information representation and reasoning space [23]. In addition to more traditional areas such as knowledge or domain model representation and

corresponding information, CE also encompasses the representation of logical inference rules, rationale (reasoning steps), assumptions, statements of truth (and certainty) and has been used in other areas such as provenance [24] and argumentation [25].

In the remainder of this section we give a number of examples of the CE language. These are shown as embedded sentences in [this style](#). All of these sentences are valid CE and therefore directly machine processable as well as being human readable.

The domain model used within CE is created through the definition of concepts, relationships and properties. These definitions are themselves written as CE conceptualise statements:

conceptualise a ~ device ~ D.
conceptualise an ~ environment variable ~ E.

These statements establish the concepts within the CE domain model enabling subsequent instances to be created using the same CE language:

there is an environment variable named 'temperature'.

A slightly more advanced example would be:

conceptualise a ~ controlling thing ~ C that
is a device and
~ can control ~ the environment variable E.

This defines “controlling thing” as a sub-concept of “device” and that it can have a “can control” relationship with an “environment variable”. This therefore allows statements such as:

there is a controlling thing named 'thermostat' that
can control the environment variable
'temperature'.

In the latter conceptualise statement, “can control” is an example of a CE verb singular relationship. Functional noun relationships can also be asserted:

conceptualise a ~ device ~ D that
has the value E as ~ enabled ~.

These two types of relationship construct allow a concept and its properties to be richly defined in CE whilst maintaining a strict subset of grammar. The use of verb singular and functional noun forms of properties provides a simple, but effective, mechanism to enable the conceptual model designer to use a language that is more natural and appealing to the human agents in the system.

The “is a” relationship used within conceptualise sentences defines inheritance of concepts, with multiple inheritance from any number of parents being a key requirement. It also allows any instance to be asserted as any number of concurrent concepts; an essential tool when attempting to capture and convey different contexts for the same information.

Whilst the examples given above are deliberately simplistic the same simple language constructs can be used to develop rich models and associated knowledge bases. The CE language has been successfully used in a wide range of example applications [26]. CE has been shown working with a reasonable number of concepts, relationships, queries and rules and has been used to model and interact with complex real-world environments with a high level of coverage and practical expressivity being achieved.

In our previous research into the application of the CE language we have observed that by gradually building up an operational model of a given environment, it is possible to iteratively define rich and complex semantic models in an “almost-NL” form that appeals to non-specialist domain users. For example, if the concept “device” was extended to include location information, the following query could be used to identify all devices of a particular type within a particular location:

for which D is it true that
(the device D is located in the room V) and
(the device D can measure
the environment variable 'temperature') and
(the value V = 'kitchen').

Note that we do not expect casual users to write CE queries of this complexity; the later conversational interaction section will show how users can do this in a more natural form.

The model can be extended with rules that can be used to automatically infer new facts within the domain. Whenever such facts are inferred the CE language is able to capture rationale for why a particular fact is held to be true:

the room 'kitchen'
is able to measure
the environment variable 'temperature' and
is able to control
the environment variable 'temperature'
because
the thermometer 't1'
is located in the room 'kitchen' and
can measure
the environment variable 'temperature' and
the radiator valve 'v1'
is located in the room 'kitchen' and
can control
the environment variable 'temperature'.

From these basic examples you can see how the CE language can be used to model the basic concepts and properties within a given domain (such as an operating environment for IoT devices). Through assertion of corresponding instance data and the use of queries and rules it is possible to define the specific details of any given environment. It should also be clear to the reader that whilst human-readable the core CE language is quite technical and does not yet meet the aspiration of a language that would appeal to everyday casual users. The language itself can be improved, and as reported in earlier research there is the ability to build incrementally usable layers of language on top of the CE core language [27]. However, in addition to all of these potential advances in the core language there is also a key innovation that has been recently developed, which is to build a rich conversational protocol on top of the CE language [28]. This provides a mechanism whereby casual users can engage in conversation with a CE knowledge base using their own NL in a manner similar to human-human conversation.

B. Conversational Interaction

To enable a conversational form of CE, earlier research [29] has identified a requirement for a number of core interaction types based on speech-act theory:

- 1) A confirm interaction allows a NL message, typically from a human user, to be provided, which is then refined through interaction to an acceptable CE representation. This is useful for a human user who is perhaps not fully trained on the CE grammar. Through multiple such interactions, their experience builds and such interactions become shorter.
- 2) An ask/tell interaction allows a query to be made of the domain model and a well-formulated CE response given.
- 3) A gist/expand interaction enables the CE agent to provide a summary form (“gist”) of a piece of CE, possibly adapted to a more digestible NL form. Such a gist can be expanded to give the underlying CE representation.
- 4) A why interaction allows an agent in receipt of CE statements to obtain rationale for the information provided.

This “conversational layer” is built within the core CE environment and is defined using the CE language. Within the CE model, these interactions are modeled as sub-types of the card concept.

conceptualise a ~ card ~ C that is an entity and
 has the timestamp T as ~ timestamp ~ and
 has the value V as ~ content ~ and
 ~ is to ~ the agent A and
 ~ is from ~ the agent B and
 ~ is in reply to ~ the card C.

The concept of an agent is introduced to represent the different parties in a conversation. This model provides a framework for such agents to interact by CE statements. By developing a conversational protocol using the CE language it enables the same language to be used for the domain in question (e.g., IoT devices in the home), as well as the act of communication. This means that agents with different operational domains can still communicate using a standard conversational model, so even if they cannot decode the items being discussed they are at least able to participate in the conversation. This idea is central to the proposed approach for conversationally enabled human and machine agents in an IoT context described in this paper.

C. Agent and ce-store interaction

In our ongoing experiments using the CE language we are able to define models, build knowledge bases, build machine agents and enable conversational interaction between them using some key components, which we will refer to here as ce-store. The Java-based implementation of the full ce-store [30] is publically available from github and an additional javascript-based version [31] is also available, specifically engineered to enable operation at the edge of the network, i.e., in a mobile browser environment.

For example, the domain model shown earlier in this paper is created through CE, (including the concepts, relationships and instances) and held within an instance of the ce-store, also referred to as a CE knowledge base. This store can either be maintained at a central point in the architecture, or distributed across systems through a federated set of separate ce-store instances. A centralized store provides a more straightforward system to maintain and ensures a single, shared model. Distributing the store allows for more localized processing to

be done by the agents without having to interact with the system as a whole. Distributing the store also enables different agents to have different models, and for models to be rapidly extended “in the field” for only those agents which require those changes.

The choice of agent architecture influences how the store should be structured. When considering the types of conversation a chat-room for the home may need to support, there are two possible approaches.

- 1) The human user interacts with a single agent in the role as a concierge for the home. This agent uses the CE knowledge base to maintain a complete situational awareness of the devices in the home and is able to communicate with them directly (see Figure 3). Interactions between concierge and devices do not use CE; only the concierge has a CE knowledge base.
- 2) The human user interacts with each device, or set of devices, individually. There may still be an agent in a concierge style role, but conversations can be directed at individual devices of interest as required (see Figure 4). Here, the concierge and all devices can communicate using CE and all have their own CE knowledge bases.




Figure 3: The human user interacts (via CE) only with the concierge

Whilst the former would be sufficient to enable purely human-machine interaction, one of the goals of this work is to enable the human to passively observe the interaction of the devices in the home in order to help the human gain awareness of how the system is behaving. This will better enable the human user to see normal behavior over time and therefore prepare them for understanding anomalous situations when they arise.



Figure 4: The human user can interact (via CE) directly with all devices and with devices via the concierge

As such, the latter approach is more suited for these purposes, perhaps with a concierge agent who is additionally

maintaining the overall situation awareness from a machine-processing perspective.

D. Modelling the Conversation

In our proposed conversational homes setting there are a number of styles of interaction a human may wish to have with the devices in their home. This section considers four such styles and how they can be handled within a CE environment.

1) *Direct question/answer exchanges*: This is where a user makes a direct query as to the current state of the environment or one of the devices therein. For example: “What is the temperature in the kitchen?”

Through the existing conversational protocol and embedded simple contextual NL processing a machine agent is able to break down such a statement to recognize its intent. By parsing each word in turn and finding matching terms with the ce-store it can establish:

- it is a question regarding a current state (“What is ...”)
- it is regarding the temperature environment variable instance
- it is regarding the kitchen room instance

At this point, the machine agent has sufficient information to query the ce-store to identify what devices in the model are in the right location and capable of measuring the required variable. If such a device exists, it can be queried for the value and reported back to the user. Otherwise, a suitable message can be returned to indicate the question cannot be answered, ideally conveying some indication of why not.

If the question is ambiguous, for example by omitting a location, the agent can prompt the user for the missing information. The concept of ambiguity for this kind of question is also captured in CE, for example by stating that for such an environment variable a location must be specified, perhaps even with a default location that can be assumed. With this knowledge available in CE the agent is able to determine that extra information is still required and can request this from the user as part of the conversation. The agent maintains information regarding the state of the conversation such that prompts can be made without requiring the user to repeat their entire question with the additional information included. By using the conversational protocol on top of the core CE language the human user and the device are able to converse in NL, for example:

User: *What is the temperature?*

Agent: *Where? I can tell you about the kitchen, the hall and the master bedroom.*

User: *The kitchen.*

Agent: *The temperature in the kitchen is 22C*

Other simple question types can be handled in this way, such as “where is...”.

2) *Questions that require a rationale as response*: This is where a user requires an explanation for a current state of the system “Why is the kitchen cold?”

As with a direct question, an agent can parse the question to identify:

- it is a question asking for a rationale (“Why is ...”)

- it has a subject of kitchen
- it has a state of cold that, through the CE model, is understood to be an expression of the temperature environment variable.

To be able to provide a response, the model supports the ability to identify what can affect the given environment variable. With that information it can examine the current state of the system to see what can account for the described state. For example, “the window is open” or “the thermostat is set to 16C”.

3) *An explicit request to change a particular state*: This is where a user, or a machine agent, makes an explicit request for a device to take an action “Turn up the thermostat in the kitchen”

To identify this type of statement, the model maintains a set of actions that can be taken and to what devices they can be applied. By incrementally matching the words of the statement against the list of known actions, a match, if it exists, can be identified. Further parsing of the statement can identify a target for the action.

conceptualise an ~ action ~ A that
~ is reversed by ~ the action B and
~ can affect ~ the controlling thing M.

if (the action A is reversed by the action B)
then (the action B is reversed by the action A).

This demonstrates the ability to define a rule. These are logic constructs with premises and conclusions that get evaluated by the ce-store against each new fact added. Where a match in the premises is found, new facts are generated using the conclusions (with corresponding rationale). In this simple case it allows two-way relationships to be established without having to explicitly define the relationship in both directions.

there is an action named 'turn on'.
there is an action named 'turn off'.
the action 'turn on' is reversed by the action 'turn off'.

When a device receives an action, the trigger concept can be used to chain further sequences of actions that should occur. For example, when applied to a thermostat, the action “turn up” should trigger the action “turn on” to be applied to the boiler.

there is a trigger named 'tr1' that
has 'turn up' as action and
has 'boiler' as target device and
has 'turn on' as target action.
the thermostat 'ts1' will respond to the trigger 'tr1'.

There is a natural point of contact here, with the popular ‘If This Then That’ framework (IFTTT) [32], specifically in that the use of conversational interactions could provide a nice way to implement IFTTT functionality. In future work we may consider the extent to which CE could be applied in IFTTT scenarios, and used to support a user-friendly form of programming for real-world objects, devices and situations.

4) *An implicit desire to change a state*: The styles considered so far have been explicit in their intent. There is another form whereby a statement is made that states a fact, but also implies a desire for an action to be taken.

This relies on Grice’s Maxim of Relevance [33]. In the context of a conversation with the devices in a house, a

statement such as "I am cold" should be taken as a desire for it to be warmer. The underlying information that can allow this Gricean inference to be implemented by machine agents using a simple algorithm is shown below:

there is a physical state named 'cold' that
is an expression of
the environment variable 'temperature' and
has 'warmer' as desired state.

there is a desired state named 'warmer' that
has 'temperature' as target and
has 'increase' as effect.

Once the intention of the statement has been identified, the store can be queried to find any actions that satisfy the requirement. These actions can then be offered as possible responses to the statement, or possibly automatically enacted.

Through these four simple dialogue examples we have demonstrated that through the use of a CE knowledge base and a set of machine agents using the conversational protocol a human user could carry out basic interactions with the devices in their home (human-machine). We have also shown how those devices convey key information back to the user, or ask follow on questions to elicit additional information (machine-human). These same interactions using the same CE language can be used to enable direct communications between machine agents regardless of human involvement (machine-machine). Whilst we have not explicitly demonstrated human-human communication it is clear that this can easily be supported within a system such as this, for example, by enabling different human users within the home to use the same chat environment to converse with each other directly and then easily direct their questions or actions to machine agents when needed.

It is the use of this common human-readable CE language that enables the passive observation of system state and agent communications at any time without development of special tooling to convert from machine specific representation formats to something that human users can directly read. The CE language enables machine or human users to change or extend the conceptual models against which the system is operating as well as allowing them to define new knowledge, queries or rules.

Whilst it would be possible to demonstrate the same capabilities using more traditional Semantic Web languages they would be aimed at machine processability rather than human consumability and would therefore require additional components to be developed to allow conversational interaction and the inclusion of the human users in the conversation.

IV. EVALUATION

As set out in the introduction, our hypothesis is that CNL can enable machine-machine, machine-human, human-machine and human-human interaction in a dynamic environment. The previous section has given illustrative examples of how we envisage the approach working in a range of use cases. Through a series of experiments, we are building an evidence base to show the feasibility and effectiveness of the approach, in two respects: (i) that humans without any significant degree of training are able to engage in dialogues using a combination of NL and CNL; and (ii) that the approach supports environments that can rapidly evolve to accommodate new devices and capabilities as they are encountered.

To gather evidence for (i), we have to date run a series of trials in controlled conditions, focusing on the proposition that users with little or no training in the use of CE can productively interact with CE-enabled agents. We reported the results of the first of these studies in [29]. Twenty participants (undergraduate students) were assigned a task of describing scenes depicted in photographs using NL, and given feedback in the form of CE statements generated via NLP by a software agent. The agent had been constructed rapidly to perform simple bag-of-words NLP with a lexicon provided by having four independent people provide scene descriptions in advance of the experiment. The results were promising; from 137 NL inputs submitted by the 20 subjects, with a median of one sentence for each input, a median of two CE elements was obtained by NLP for each input. In other words, with no prior training in the use of CE or prior knowledge of the domain model constructed for the scenes, users were able to communicate two usable CE elements (typically an identified instance and a relationship) per single-sentence NL input.

The ability of the CE agent to extract meaningful elements from the user's input and confirm these in CE form was constrained by the rapid manner in which the background domain knowledge base had been constructed. In effect, the agent's limited knowledge about the world led to results that were characterized by high precision, but relatively low recall, since the agent was engineered only to be "interested" in a narrow range of things. In this respect, however, we see these results as applicable to our "conversational homes" scenarios, where the concerns of home-based devices and the affordances users expect them to provide will be similarly narrow. Further experiments are planned in settings more closely aligned with the examples in the previous section.

In terms of our requirement (ii), that the approach supports environments which can rapidly evolve to accommodate new devices and capabilities as they are encountered, we have constructed and demonstrated experimental prototypes for sensing asset selection for users' tasks, as described in [34]. Again, while these prototypes are not exactly aligned with the scenario of home automation (instead being more concerned with sensing systems such as autonomous aerial vehicles and ground systems) these experiments have shown that the CE-based approach supports the rapid addition of new knowledge. This includes not only of types of asset, but also of asset capabilities (that can be used to match assets to tasks). In many ways, the home setting is simpler than, say, an emergency response or search-and-rescue scenario, so we believe that the positive outcomes of these experiments are translatable into the domestic context.

An arguable difference between the home versus emergency response or search-and-rescue settings is the degree of training that a user can reasonably be expected to have obtained in the use of the available devices. In the home setting, this must always be minimal. In the other setting, however, minimal training is still desirable, since users should not necessarily be experts in the operation of sensing systems [35]. In any case, we argue that this usability point is addressed under (i) above. Also, in many cases, the addition of knowledge about new devices and their capabilities will typically be provided by the originators of the devices rather than end-users, though our approach does not preclude a "power" user from providing additional knowledge to their local environment.

V. CONCLUSION

In this paper we have explored the use of conversations between humans and machines, motivated by a desire for “beautiful seams”. We assert that this approach could enable better understanding of complex system such as a set of IoT devices in a home. In this paper, we have shown how semantic representations can be used in a human-friendly format through the use of a CNL technology known as ITA CE. Through the use of a conversational protocol built on top of the core CE language we show how human and machine agents are able to communicate using this single language. Examples of the CE language are provided throughout the paper showing how different concepts can be constructed and the subsequent data for the knowledge base can be provided in the same CE language. Through a set of four typical types of interaction we show how human users can interact with the devices in such an environment, and we note that whilst we have focused these four examples on a human-machine interaction, the exact same approach applies to machine-machine as well. Some additional discussion around what machine-human and human-human forms would look like is mentioned. Future work may include conducting experiments in the conversational home setting, aiming to replicate the results from our earlier work where human users without training were able to use the conversational protocol and the CE language to communicate key features within the domain of interest.

ACKNOWLEDGMENT

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-16-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copy-right notation hereon.

REFERENCES

- [1] R. Dale, “The return of the chatbots,” *Natural Language Engineering*, vol. 22, no. 5, pp. 811–817, 2016.
- [2] M. Witbrock and L. Bradeško, “Conversational computation,” in *Handbook of Human Computation*. Springer, 2013, pp. 531–543.
- [3] J. Lester, K. Branting, and B. Mott, “Conversational agents,” *The Practical Handbook of Internet Computing*, pp. 220–240, 2004.
- [4] N. O’Leary. (2014) Conversational iot. [Online]. Available: <http://knollearly.net/2014/12/04/a-conversational-internet-of-things-thingmonk-talk/>
- [5] A. Stanford-Clark. (2008) Redjets on twitter. [Online]. Available: <https://twitter.com/redjets>
- [6] T. Armitage. (2008) Tower bridge on twitter. [Online]. Available: https://twitter.com/twrbrdg_itself
- [7] (2008) Mars curiosity on twitter. [Online]. Available: <https://twitter.com/MarsCuriosity>
- [8] (2010) Philae lander on twitter. [Online]. Available: <https://twitter.com/Philae2014>
- [9] (2011) Rosetta probe on twitter. [Online]. Available: https://twitter.com/ESA_Rosetta
- [10] E. Ferrara, O. Varol, C. Davis, F. Menczer, and A. Flammini, “The rise of social bots,” *Communications of the ACM*, vol. 59, no. 7, pp. 96–104, 2016.
- [11] L. Atzori, A. Iera, and G. Morabito, “The internet of things: A survey,” *Computer networks*, vol. 54, no. 15, pp. 2787–2805, 2010.
- [12] M. Weisser, “The computer for the twenty-first century,” *Scientific American*, vol. 265, no. 3, pp. 94–104, 1991.
- [13] M. Chalmers, “Seamful design and ubicomp infrastructure,” in *Proceedings of Ubicomp 2003 Workshop at the Crossroads: The Interaction of HCI and Systems Issues in Ubicomp*. Citeseer, 2003.
- [14] T. Coates. (2014) Interacting with a world of connected objects. [Online]. Available: <https://medium.com/product-club/interacting-with-a-world-of-connected-objects-875b4a099099#nd00bbs5n>
- [15] Hypercat. [Online]. Available: <http://hypercat.io>
- [16] M. Bates and R. M. Weischedel, *Challenges in natural language processing*. Cambridge University Press, 2006.
- [17] T. Kuhn, “A survey and classification of controlled natural languages,” *Computational Linguistics*, vol. 40, no. 1, pp. 121–170, 2014.
- [18] R. Schwitter, “Controlled natural languages for knowledge representation,” in *Proceedings of the 23rd International Conference on Computational Linguistics: Posters*. Association for Computational Linguistics, 2010, pp. 1113–1121.
- [19] D. Mott, “Summary of ita controlled english,” *ITA Technical Paper*, <http://www.usukita.org>, 2010.
- [20] A. Preece and W. R. Sieck, “The international technology alliance in network and information sciences,” *IEEE Intelligent Systems*, vol. 22, no. 5, 2007.
- [21] D. Mott, C. Giammanco, M. C. Dorneich, J. Patel, and D. Braines, “Hybrid rationale and controlled natural language for shared understanding,” *Proc. 6th Knowledge Systems for Coalition Operations*, 2010.
- [22] T. Klapiscaak, J. Ibbotson, D. Mott, D. Braines, and J. Patel, “An interoperable framework for distributed coalition planning: The collaborative planning model,” *Proc. 7th Knowledge Systems for Coalition Operations*, 2012.
- [23] D. Braines, D. Mott, S. Laws, G. de Mel, and T. Pham, “Controlled english to facilitate human/machine analytical processing,” *SPIE Defense, Security, and Sensing*, pp. 875 808–875 808, 2013.
- [24] J. Ibbotson, D. Braines, D. Mott, S. Arunkumar, and M. Srivatsa, “Documenting provenance with a controlled natural language,” in *Annual Conference of the International Technology Alliance (ACITA)*, 2012.
- [25] F. Cerutti, D. Mott, D. Braines, T. J. Norman, N. Oren, and S. Pipes, “Reasoning under uncertainty in controlled english: an argumentation-based perspective,” *AFM*, 2014.
- [26] D. Braines, J. Ibbotson, D. Shaw, and A. Preece, “Building a living database for human-machine intelligence analysis,” in *Information Fusion (Fusion), 2015 18th International Conference on*. IEEE, 2015, pp. 1977–1984.
- [27] D. Mott and J. Hendler, “Layered controlled natural languages,” in *3rd Annual Conference of the International Technology Alliance (ACITA)*, 2009.
- [28] A. Preece, D. Braines, D. Pizzocaro, and C. Parizas, “Human-machine conversations to support multi-agency missions,” *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 18, no. 1, pp. 75–84, 2014.
- [29] A. Preece, C. Gwilliams, C. Parizas, D. Pizzocaro, J. Z. Bakdash, and D. Braines, “Conversational sensing,” in *SPIE Sensing Technology+ Applications*. International Society for Optics and Photonics, 2014, pp. 91 220I–91 220I.
- [30] D. Braines. (2015) Ita controlled english store (ce-store). [Online]. Available: <https://github.com/ce-store>
- [31] W. Webberley. (2016) Cenode.js. [Online]. Available: <http://cenode.io/>
- [32] If this then that. [Online]. Available: <https://ifttt.com/>
- [33] H. P. Grice, “Logic and conversation,” *1975*, pp. 41–58, 1975.
- [34] A. Preece, D. Pizzocaro, D. Braines, and D. Mott, “Tasking and sharing sensing assets using controlled natural language,” in *SPIE Defense, Security, and Sensing*. International Society for Optics and Photonics, 2012, pp. 838 905–838 905.
- [35] A. Preece, T. Norman, G. de Mel, D. Pizzocaro, M. Sensoy, and T. Pham, “Agilely assigning sensing assets to mission tasks in a coalition context,” *IEEE Intelligent Systems*, vol. 28, no. 1, pp. 57–63, 2013.