

Assembling the jigsaw: How multiple open standards are synergistically combined in the HALEF multimodal dialog system

Vikram Ramanarayanan[†], David Suendermann-Oeft[‡], Patrick Lange[†],
Robert Mundkowsky[‡], Alexei V. Ivanov[†], Zhou Yu*, Yao Qian[†] & Keelan Evanini[‡]

As dialog systems become increasingly multimodal and distributed in nature with advances in technology and computing power, they become that much more complicated to design and implement. However, open industry and W3C standards provide a silver lining here, allowing the distributed design of different components that are nonetheless compliant with each other. In this chapter we examine how an open-source, modular, multimodal dialog system – HALEF – can be seamlessly assembled, much like a jigsaw puzzle, by putting together multiple distributed components that are compliant with the W3C recommendations or other open industry standards. We highlight the specific standards that HALEF currently uses along with a perspective on other useful standards that could be included in the future. HALEF has an open codebase to encourage progressive community contribution and a common standard testbed for multimodal dialog system development and benchmarking.

1 Introduction

Dialog systems nowadays are becoming increasingly multimodal. In other words, dialog applications, which started off mostly based on voice and text [15], have increasingly started to encompass other input-output (I/O) modalities such as video [3], gesture [18, 3], electronic ink [11, 10], avatars or virtual agents [6, 27, 26], and even embodied agents such as robots [7, 29], among others. While the integration of such technologies provides a more immersive and natural experience for the users and enables an analysis of their non-verbal behaviors, it also makes the design of such multimodal dialog systems more complicated. This is because, among other things, one needs to ensure a seamless user experience without any reduction in quality of service – this includes issues such as latency, accuracy and sensitivity – while transporting data between each of these multimodal (and possibly disparate) I/O endpoints and the dialog system. In addition, dialog systems consist of multiple subsystems; for example, automatic speech recognizers (ASRs), spoken language understanding (SLU) modules, dialog managers (DMs) and speech synthesizers,

[†] Educational Testing Service (ETS) R&D, San Francisco, CA

[‡] Educational Testing Service (ETS) R&D, Princeton, NJ

* Carnegie Mellon University, Pittsburgh, PA

Email: <vramanarayanan@ets.org>

among others, interacting synergistically and often in real-time. Each of these subsystems is complex and brings with it design challenges and open research questions in its own right. As a result, development of such multi-component systems that are capable of handling a large number of calls are typically done by large industrial companies and a handful of academic research labs since they require individual maintenance of multiple individual subsystems [5]. In such scenarios, it is essential to have industry-standard protocols and specification languages that ensure interoperability and compatibility of different services, irrespective of who designed them or how they were implemented. Designing systems that adhere to such standards also allows generalization and accessibility of contributions from a large number of developers across the globe.

The popularity of commercial telephony-based spoken dialog systems – also known as interactive voice response (IVR) systems – especially in automating customer service transactions in the late 1990s, drove industry developers to start working on standards for such systems [19]. As a core component of an IVR, the voice browser, essentially responsible for interpreting the dialog flow while simultaneously orchestrating all the necessary resources such as speech recognition, synthesis and telephony, was one of the early components subject to standardization resulting in the VoiceXML standard dating back to 1999¹ (see Section 4.1.1 for more details on VoiceXML). Since the vast majority of authors responsible for creating standards such as VoiceXML come from the industry, most implementations of spoken dialog systems adhering to these standards are commercial, proprietary, and closed-source applications. Examples of voice browser implementations include:

- Voxeo Prophecy²
- TellMe Studio³
- Plum DEV⁴
- Cisco Unified Customer Voice Portal⁵
- Avaya Voice Portal⁶

In addition to over 20 commercial closed-source voice browsers⁷, we are aware of a single open-source implementation that has been actively developed over the past few years:

- *JVoiceXML*⁸.

We adopted this voice browser for the creation of the multimodal spoken dialog system HALEF (Help Assistant–Language-Enabled and Free), which serves as an example of a standards-based architecture in this chapter.

¹ <http://www.w3.org/TR/2000/NOTE-voicexml-20000505>

² <https://voxeo.com/prophecy/>

³ <https://studio.tellme.com/>

⁴ <http://www.plumvoice.com/products/plum-d-e-v/>

⁵ <http://www.cisco.com/c/en/us/products/customer-collaboration/unified-customer-voice-portal>

⁶ <https://support.avaya.com/products/P0979/voice-portal>

⁷ Find a comprehensive list at <https://www.w3.org/Voice/voice-implementations.html>

⁸ <https://github.com/JVoiceXML/JVoiceXML>

Note that in addition to industrial implementations of spoken and multimodal dialog systems, there exists an active academic community engaging in research on such systems. Prominent examples include:

- CMU’s Olympus [4]
- Alex⁹, by the Charles University in Prague [12]
- InproTK¹⁰, an incremental spoken dialog system
- OpenDial¹¹
- the Virtual Human Toolkit [9]
- Metalogue¹², a multimodal dialog system
- IrisTK¹³, a multimodal dialog system

Many of these examples, along with other (multimodal) dialog systems developed by the academic community, are built around very specific research objectives. For example, Metalogue provides a multimodal agent with metacognitive capabilities; InproTK was developed mainly for investigating the impact of incremental speech processing on the naturalness of human-machine conversations; OpenDial allows one to compare the traditional MDP/POMDP¹⁴ dialog management paradigm with structured probabilistic modelling [14]. Due to their particular foci, they often use special architectures, interfaces, and languages paying little attention to existing speech and multimodal standards (e.g. see the discussions in [2]). For example, none of the above research systems implements VoiceXML, MRCPv2, or EMMA (see Section 4 for more details on these standards).

In this chapter, we describe a system that was designed to bridge the gap between the industrial demand for standardization and the openness, community engagement, and extensibility required by the scientific community. This system, HALEF, is an open-source cloud-based multimodal dialog system that can be used with different plug-and-play back-end application modules [22, 25, 30]. In the following sections, we will first describe the overall architecture of HALEF (Section 2) including its operational flow explaining how multimodal interactions are carried out (in Section 3). We will then review major components of multimodal dialog systems that have previously been subject to intensive standardization activity by the international community and discuss to what extent these standards are reflected in the HALEF framework. These include:

- standards for **dialog specification** describing system prompts, use of speech recognition and interpretation, telephony functions, routing logic, etc. (primarily VoiceXML), see Section 4.1.1 (also see the chapter on the topic in this book [1]);
- standards controlling properties of the **speech recognizer**, primarily grammars, statistical language models, and semantic interpretation (e.g., JSGF, SRGS, ARPA, WFST, SISR), see Section 4.1.2;
- standards controlling properties of the **speech synthesizer** (primarily SSML);

⁹ <https://github.com/UFAL-DSG/alex>

¹⁰ <https://bitbucket.org/inpro/inprotk>

¹¹ <http://www.opardial-toolkit.net>

¹² <http://www.metalogue.eu>

¹³ <http://www.iristk.net>

¹⁴ Partially Observable Markov Decision Processes

- standards controlling the **communication** between the components of the multimodal dialog system (SIP, MRCPv2, WebRTC, EMMA), see Section 4.2;
- standards describing the **dialog flow** and how **modalities** interact (SCXML, EMMA), see Section 5.

2 The HALEF Dialog System

The multimodal HALEF framework [22, 25, 30] is composed of the following distributed open-source modules (see Figure 1 for a schematic overview):

- Telephony servers – Asterisk [16] and Freeswitch [17] – that are compatible with SIP (Session Initiation Protocol), PSTN (Public Switched Telephone Network) and WebRTC (Web Real-Time Communications) standards, and include support for voice and video communication.
- A voice browser – JVoiceXML [23] – that is compatible with VoiceXML 2.1, can process SIP traffic, via a voice browser interface called Zanzibar [21] and incorporates support for multiple grammar standards such as JSGF (Java Speech Grammar Format), ARPA (Advanced Research Projects Agency), and WFST (Weighted Finite State Transducer), which are described in Section 4.1.2.
- An MRCPv2 (Media Resource Control Protocol Version 2) speech server – which allows the voice browser to control media processing resources such as speech recorders, speech recognizers or speech synthesizers over the network. It relies on other protocols such as SIP for session handling, RTP (Real-time Transport Protocol) for media streaming and SDP (Session Description Protocol) to allow the exchange of other capabilities such as supported codecs over the network. HALEF supports multiple speech recognizers (Sphinx [13], Kaldi [20]) and synthesizers (Mary [24], Festival [28]).
- A webserver – Apache Tomcat¹⁵ that can host web applications that serve dynamic VoiceXML pages, web services, as well as media libraries containing grammars and audio files.
- OpenVXML, a voice application authoring suite that generates dynamic web applications that can be housed on the web server (also see Section 4.1.1).
- A MySQL¹⁶ database server for storing call log information. All modules in HALEF connect to the database and write their log messages to it. We then post-process this information with stored procedures into easily accessible views.
- A custom-developed, open-source Speech Transcription, Annotation and Rating (STAR) portal that we implemented using PHP and the JavaScript framework jQuery. The portal allows one to analyze, listen to (or watch) full-call (video) recordings, transcribe them, rate them on a variety of dimensions such as caller experience and latency, and perform various semantic annotation tasks required to train automatic speech recognition and spoken language understanding modules.

¹⁵ <http://tomcat.apache.org/>

¹⁶ <https://www.mysql.com/>

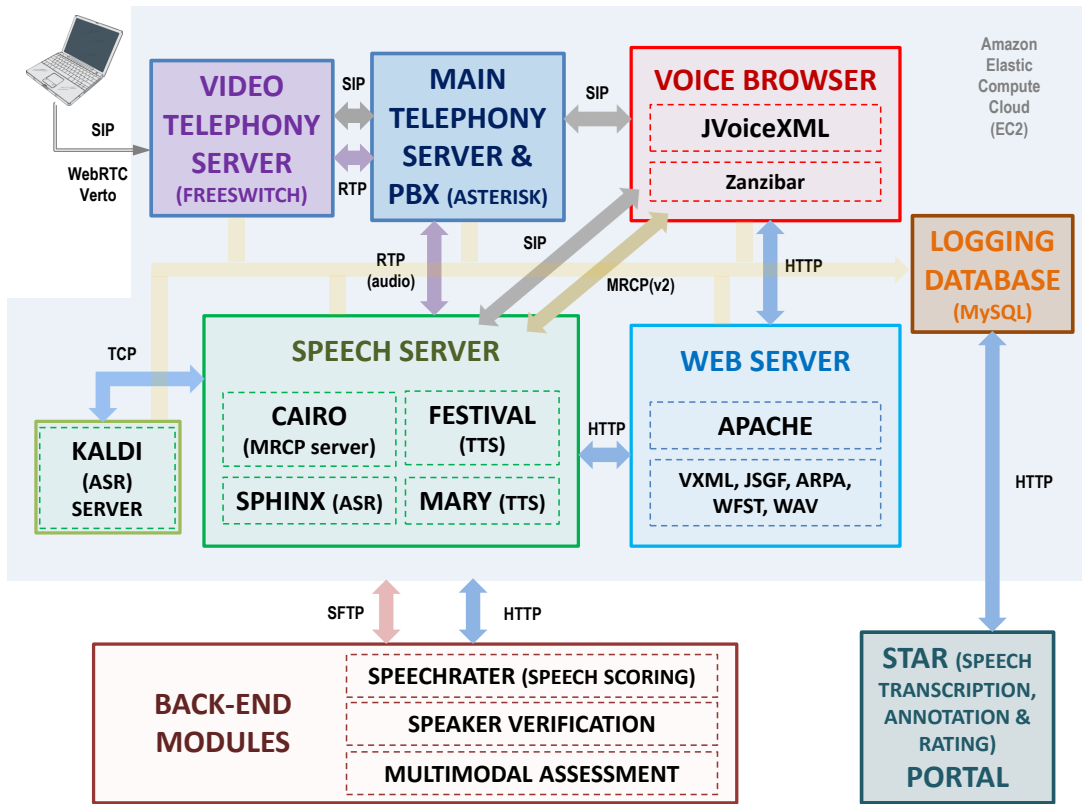


Fig. 1 System architecture of the HALEF spoken dialog system depicting the various modular open-source components as well as W3C standard protocols that are employed.

We will illustrate the basic architecture and components of the HALEF spoken dialog system using an example application that is currently deployed in the educational domain. Finally we will conclude with a discussion of ongoing and future research and development into the system, including potential support for additional W3C standards such as EMMA (Extensible Multimodal Annotation), SSML (Speech Synthesis Markup Language), EmotionML (Emotion Markup Language), and SCXML (State Chart XML).

3 Operational flow schematic

In this section we describe how video and audio data flow to/from the multimodal HALEF system. In case of regular PSTN telephony, users call into a phone num-

ber which connects them to the telephony server in the cloud where they need to provide an *extension* to connect to (different extensions are associated with different dialog system instances that in turn have different task content). Alternatively, users can use softphones (or SIP phones) to connect directly to the IP address of the cloud-based telephony server using the extension. Even more convenient is the use of a web application to call directly out of a web browser. Here, the only information required by the user is the URL of the website containing the connection configuration (which includes the telephony server IP address and the extension). This web application is written in HTML, CSS, and Javascript. The Media Capture and Streams API¹⁷ enables access to the computer's audio and video input devices via the web browser. WebRTC¹⁸ is then used via a Javascript implementation to send video and audio to FreeSWITCH and receive audio back from FreeSWITCH. When the call comes in from the user, HALEF starts the dialog with an audio prompt that flows out of the HALEF system via Asterisk over SIP/RTP to FreeSWITCH. FreeSWITCH then sends the audio to the web browser via WebRTC. The user then gives a response to the system that flows through WebRTC to FreeSWITCH and then through SIP/RTP to Asterisk. During the teleconference, the user's video and audio interactions are continuously streamed and recorded.

Once the Asterisk server receives the call, it sends a notification to the voice browser to fetch the VXML code from the web server. The voice browser in turn identifies the resources that the speech server will need to prepare for this application. It then notifies the MRCPv2 server and starts sessions and channels for all required resources including the provisioning of speech recognition grammars. Finally, the speech server sends a SIP response back to the voice browser and Asterisk to confirm session initiation. Completion of this process successfully establishes a communication channel between the user and HALEF's components. Once the session is established, Asterisk streams audio via RTP to the speech server. When the caller starts speaking, the Sphinx engine's voice activity detector fires and identifies voiced portions of the speech; then, the speech is sent to the ASR engine (HALEF supports both Kaldi and Sphinx) which starts the decoding process. When the voice activity detector finds that the caller has finished speaking, the recognition result is sent back to the voice browser, which processes it and sends this answer to the spoken language understanding module. The output of the natural language understanding module is subsequently sent to the dialog manager which evaluates and generates VXML code with the final response to be spoken out by the speech synthesizer (either Festival or Mary). The voice browser then interprets this VXML code and sends a synthesis request to the speech server with the response. Festival/Mary synthesizes the response and passes the result back via RTP to Asterisk, which forwards the audio signal to the user. At the same time, Cairo sends a confirmation signal to the voice browser. After receiving this signal, the voice browser sends a cleanup request to close all open channels and resources. This ends the SIP session with Asterisk, which finally triggers Asterisk to send an end-of-call signal to the user.

There are other endpoints that are supported or likely can be supported by HALEF. An endpoint is defined as a device at the edge of the network (e.g. a

¹⁷ <https://www.w3.org/TR/mediacapture-streams>

¹⁸ <http://www.w3.org/TR/webrtc/>

telephone or a soft phone). Note that HALEF also natively supports audio-only dialogs with PSTN (public switched telephone network) or soft phone endpoints (that, for example, can use PSTN/SIP proxies such as ipKall¹⁹. We have also used SIP over WebRTC, and SIP/WebRTC clients such as sipml5²⁰, jssip²¹, etc. to connect to HALEF directly through Asterisk as well as via webrtc2sip²² to Asterisk.

4 Standards used in HALEF

The following section examines in more detail how different specific industry standard specifications are synergistically combined within the HALEF multimodal dialog framework. Since HALEF is primarily a spoken dialog system, we first examine the key voice standards used in its operation. We then describe the various communication standards used to transport voice and video data across different components of the dialog system.

4.1 Voice Standards

4.1.1 VoiceXML

The origins of VoiceXML²³ began in 1995 as an XML-based dialog design language intended to simplify the speech recognition application development process within an AT&T project called Phone Markup Language (PML). VoiceXML or VXML was designed for creating audio dialogs that feature synthesized speech, digitized audio, recognition of spoken and DTMF key input, recording of spoken input, telephony, and mixed initiative conversations. It was conceived to integrate the advantages of web-based development and content delivery into interactive voice response applications. The code listing below shows an example VXML page as used by HALEF. This example VXML page illustrates how various system parameters can be specified, such as the timeout value of 3 seconds specified in the `timeout` variable. Also, this example shows several of the components required for the interactive conversation, such as the system prompt (a prerecorded audio file, in this case) specified in the `<prompt>` element and the grammar file (see Section 4.1.2) that controls which user utterances can be recognized by the ASR system.

```
<vxml version="2.1">
<form id="InputRequestForm" scope="document">
  <field name="A_try_peanuts">
    <property name="bargein" value="true"/>
  </field>
</form>
```

¹⁹ <http://www.ipkall.com/>

²⁰ <https://www.doubango.org/sipml5/>

²¹ <http://www.jssip.net/>

²² <http://webrtc2sip.org/>

²³ <http://www.w3.org/TR/voicexml20/>

```

    <property name="timeout" value="3s"/>
    <property name="confidencelevel" value="0.5"/>
    <property name="sensitivity" value="0.5"/>
    <property name="speedvsaccuracy" value="0.5"/>
    <property name="completetimeout" value="3s"/>
    <property name="incompletetimeout" value="3s"/>
    <property name="maxspechtimeout" value="10s"/>
    <property name="inputmodes" value="voice"/>
    <property name="com.telera.speechenabled" value="true"/>
    <prompt bargein="true" xml:lang="en-US">
      <audio
        src="/7703/-/resources/EPS_Builder_Voice/Default/peanuts_offer.wav"/>
    </prompt>
    <grammar mode="voice" type="application/srgs+xml"
      src="/7703/-/resources/EPS_Builder_Voice/Default/try_peanuts.gram"/>
    <filled>
    <var name="lastresult" expr="'<lastresult>'" />
    <submit
      next="/7703/-/next?Action_216121ee52ce43378ca2e014b92f71b4=success.filled"
      method="post" namelist="A_try_peanuts lastresult"/>
    </filled>
    <noinput></noinput>
    <nomatch></nomatch>
    <catch event="connection.disconnect.hangup"></catch>
  </field>
<catch event="externalmessage.cpa.machine"></catch>
<catch event="externalmessage.cpa.beep"></catch>
<catch event="externalmessage.cpa.machine"></catch>
</form>
<catch event="connection.disconnect.hangup"></catch>
</vxml>

```

However, developers of dialog applications who are not familiar with the VXML markup language may prefer to define dialog flows using a simpler, flowchart-based GUI instead of manual coding. Therefore we have integrated OpenVXML (or Open Voice XML) into the HALEF framework. OpenVXML is an open-source software package²⁴ written in Java that allows designers to author dialog workflows using an easy-to-use graphical user interface, and is available as a plugin to the Eclipse Integrated Developer Environment²⁵. OpenVXML allows designers to specify the dialog workflow as a flowchart, including details of specific grammar files to be used by the speech recognizer and text-to-speech prompts that need to be synthesized. In addition, they can insert “script” blocks of Javascript code into the workflow that can be used to perform simple processing steps, such as basic natural language understanding on the outputs of the speech recognition. The entire workflow can be exported to a Web Archive (or WAR) application, which can then be deployed on a web server running Apache Tomcat that serves Voice XML (or VXML) documents.

Figure 2 shows a simple OpenVXML dialog flow where callers are required to accept or decline an offer of food in a pragmatically appropriate manner. This example can be compared to the example VXML page shown in the earlier code listing to illustrate the differences between designing a dialog directly using VXML or

²⁴ <https://github.com/OpenMethods/OpenVXML>

²⁵ www.eclipse.org

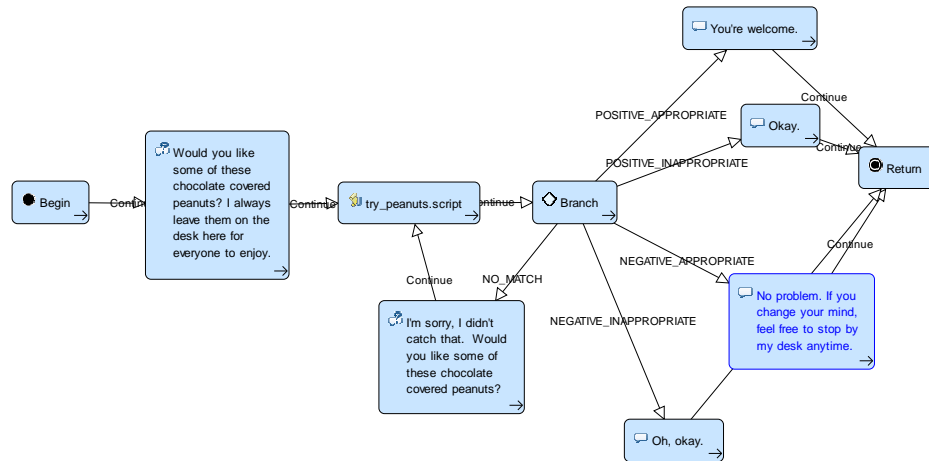


Fig. 2 Example design of a workplace pragmatics-oriented application targeted at non-native speakers of English where the caller has to accept or decline an offer of food (peanuts, in this case) in a pragmatically appropriate manner.

through the OpenVXML authoring tool. The VXML code therein corresponds to the first block in Figure 2 in which a system prompt is played (“Would you like some of these chocolate covered peanuts? ...”) By double-clicking on this block in the OpenVXML tool, the designer specifies the prompt that should be played or generated by the TTS engine (as indicated in the `<prompt>` element in the VXML page), the grammar that should be used to recognize the utterance by the ASR system (corresponding to the `<grammar>` element in the VXML page), as well as a variety of system parameters, such as the timeout variable. This GUI-based representation in OpenVXML is then translated into VXML pages automatically so that it can be deployed to the voice browser.

This item was designed to measure two primary constructs of English language proficiency: (i) task comprehension, i.e., correctly understanding the stimulus material and the questions being asked and (ii) pragmatic appropriateness, i.e., the ability to provide a response that is appropriate to the task and the communicative context. The caller dials into the system and then proceeds to answer one or more questions, which can be either be stored for later analysis (so no online recognition and natural language understanding is needed), or processed in the following manner: depending on the semantic class of the callers’ answer to each question (as determined by the output of the speech recognizer and the natural language understanding module), they are redirected to the appropriate branch of the dialog tree and the conversation continues until all such questions are answered.

4.1.2 Voice Grammar and Language Model Standards

Grammars are used by speech recognizers to determine what a speech recognizer should listen for, and so describe the utterances a user may say. This section describes the standard grammar formats that are supported by the HALEF system.

1. JGSF:

The JSpeech Grammar Format (JSGF²⁶) is a platform- and vendor-independent textual representation of grammars for use in ASR. It adopts the style and conventions of the Java Programming Language in addition to use of traditional grammar notations. For example, the following JSGF grammar accepts one of two speech recognition outputs, “yes” or “no”.

```
#JSGF V1.0;
grammar yesno;
public <yesno> = yes | no;
```

2. SRGS:

The Speech Recognition Grammar Specification (SRGS²⁷) allows the grammar syntax to be written in one of two forms – an Augmented Backus-Naur Form (ABNF) or an Extensible Markup Language (XML) form – which are semantically mappable to allow transformations between themselves. The following code snippet shows how a yes/no grammar can be defined in the ABNF format of SRGS.

```
#ABNF 1.0 UTF-8;
language en-US; //use the American English pronunciation
dictionary.
mode voice; //the input for this grammar will be spoken
words.

root $yesorno;

$yes = yes;
$no = no;
$yesorno = $yes | $no;
```

3. ARPA:

Although not a W3C recommendation, the Advanced Research Projects Agency (ARPA) format was one of the first popular ones that allowed specification of *statistical* grammars (also called language models or LMs) such as finite state automata (FSAs) or statistical n-gram models. The language model is a list of possible word sequences. Each sequence listed has an associated statistically estimated language probability tagged to it. The following listing shows an example of a yes/no ARPA grammar.

²⁶ JSGF (see <http://www.w3.org/TR/jsgf/>) is technically not a W3C standard. It is a member submission and is published as a W3C note.

²⁷ <http://www.w3.org/TR/speech-grammar/>

```

This is an example ARPA-format language model file
\data\
ngram 1=4
ngram 2=4
ngram 3=4

\1-grams:
-0.7782 </s> -0.1761
-0.3010 <s> -0.5228
-0.7782 no -0.3978
-0.7782 yes 0.0000

\2-grams:
-0.1761 </s> <s> -0.0791
-0.3978 <s> no 0.1761
-0.3978 <s> yes -0.2217
-0.1761 no </s> 0.1761

\3-grams:
-0.3010 </s> <s> yes
-0.3010 <s> no </s>
-0.3010 <s> yes </s>
-0.3010 no </s> <s>

\end\

```

4. WFST:

Speech and dialog system developers nowadays are increasingly moving to the Weighted Finite State Transducer (WFST) representation to write statistical grammars for their applications owing to its simplicity and power, even though it is not an official W3C recommendation. WFSTs are automata where each transition has an input label, an output label, and a weight. The weights can be used to represent the cost of taking a particular transition. The following shows an example of a WFST grammar (in text form) that accepts the words “yes” or “no”.

```

# arc format: src dest ilabel olabel [weight]
# final state format: state [weight]
# lines may occur in any order except initial state must
  be first line
# unspecified weights default to 0.0 (for the
  library-default Weight type)
0 1 yes yes 0.5
0 1 no no 1.5
1 2.0
EOF

```

4.2 Communication Standards

WebRTC²⁸ or Web Real-Time Communication is a free, open W3C project that provides browsers and mobile applications with Real-Time Communications (RTC) capabilities via simple APIs. It defines a set of ECMAScript APIs in WebIDL to allow media to be sent to and received from another browser or device implementing the appropriate set of real-time protocols. As explained earlier, HALEF leverages the Verto protocol implemented in the Freeswitch video telephony server that is WebRTC-based to transmit video and audio data between the user and the dialog system.

The Media Resource Control Protocol Version 2 (MRCPv2) is a standard communication protocol for speech resources (such as speech recognition engines, speech synthesis engines, etc.) across VoIP networks which is designed to allow a client device to control media processing resources on the network.

5 Other useful standards for multimodal dialog systems

There are several other useful standards that we are exploring for potential future integration into the HALEF framework. This section takes a closer look at some of these standards.

EMMA

The Extensible MultiModal Annotation (EMMA²⁹) markup language is intended for use by systems that provide semantic interpretations for a variety of inputs, including but not necessarily limited to, speech, natural language text, GUI and ink input. The language is focused on annotating single inputs from users, which may be either from a single mode or a composite input combining information from multiple modes, as opposed to information that might have been collected over multiple turns of a dialog. The language provides a set of elements and attributes that are focused on enabling annotations on user inputs and interpretations of those inputs. EMMA would be a very useful standard to integrate into the HALEF framework given the focus on multimodal dialog, and hence this is one standard we are looking to include support for in HALEF going forward.

EmotionML

Emotion Markup Language or EmotionML³⁰, as the name suggests, is “intended to be a standard specification for processing emotions in applications such as: (i)

²⁸ See <http://www.w3.org/TR/webrtc/> and <https://webrtc.org/>

²⁹ <http://www.w3.org/TR/emma>

³⁰ <https://www.w3.org/TR/emotionml/>

manual annotation of data; (ii) automatic recognition of emotion-related states from user behavior; and (iii) generation of emotion-related system behavior”. Given the importance and ubiquity of emotions in dialog interactions and the subsequent requirement for automated analysis and processing of emotional state data, developing systems that are compatible with EmotionML would extend the accessibility and generalizability of those systems.

SCXML

State Chart XML (SCXML³¹) is, according to the spec, “a general-purpose event-based state machine language that combines concepts from Call Control eXtensible Markup Language (CCXML) and Harel State Tables”. CCXML³² is “an event-based state machine language designed to support call control features in Voice Applications (including, but not limited to, VXML). The CCXML 1.0 specification defines both a state machine and event handing syntax and a standardized set of call control elements”. Harel State Tables are a state machine notation that was developed by the mathematician David Harel [8]. They offer a clean and well-thought out semantics for sophisticated constructs such as parallel states. Although we don’t require an additional state machine language as such in the current version of HALEF for smooth function, including support for this would lead to an expanded dialog functionality, allowing one to specify dialog trees as generic state machines.

6 Conclusions and Outlook

We have presented the current state of the art of the HALEF system – a fully open-source, modular, and standards-compliant spoken dialog system that can be interfaced with a number of potential back-end applications. We have illustrated the various open and W3C recommendations such as VoiceXML, WebRTC and MR-CPv2, among others, associated with different parts of the HALEF operational flow, demonstrating how these help in seamlessly assembling multiple components into a fully-functional multimodal dialog system. HALEF sourcecode can be accessed online at the following URL: <http://sourceforge.net/projects/halef/>.

There remain many exciting directions for future research and development. For instance, the current HALEF implementation allows for audio and video input from the user and can synthesize output audio, but does not support full-fledged multimodal synthesis. In the future we would like to be able to incorporate support for video and emotion generation, as well as the control of avatars and simulations. Additionally, we would like to incorporate W3C recommendations such as EMMA and EmotionML into the HALEF architecture, which would maintain standards-compliance and be useful to support such an endeavor as well.

³¹ <https://www.w3.org/TR/scxml/>

³² <https://www.w3.org/TR/ccxml/>

References

1. Baggia, P., Burnett, D., Marchand, R., Matula, V.: The role and importance of speech standards. In: *Multimodal Interaction with W3C Standards: Towards Natural User Interfaces to Everything*, p. to appear. Springer (2016)
2. Baumann, T., Buß, O., Schlangen, D.: *Inprokt in action: Open-source software for building german-speaking incremental spoken dialogue systems* (2010)
3. Bohus, D., Horvitz, E.: Facilitating multiparty dialog with gaze, gesture, and speech. In: *International Conference on Multimodal Interfaces and the Workshop on Machine Learning for Multimodal Interaction*, p. 5. ACM (2010)
4. Bohus, D., Raux, A., Harris, T., Eskenazi, M., Rudnicky, A.: Olympus: An Open-Source Framework for Conversational Spoken Language Interface Research. In: *Proc. of the HLT-NAACL*. Rochester, USA (2007)
5. Damnati, G., Béchet, F., De Mori, R.: Experiments on the france telecom 3000 voice agency corpus: academic research on an industrial spoken dialog system. In: *Proceedings of the Workshop on Bridging the Gap: Academic and Industrial Research in Dialog Technologies*, pp. 48–55. Association for Computational Linguistics (2007)
6. DeMara, R.F., Gonzalez, A.J., Jones, S., Johnson, A., Hung, V., Leon-Barth, C., Dookhoo, R.A., Leigh, J., Renambot, L.: Towards interactive training with an avatar-based human-computer interface. In: *The Interservice Industry Training, Simulation & Education Conference, ITSEC* (December 2008). Citeseer (2008)
7. Gorostiza, J.F., Barber, R., Khamis, A.M., Pacheco, M., Rivas, R., Corrales, A., Delgado, E., Salichs, M.A.: Multimodal human-robot interaction framework for a personal robot. In: *Robot and Human Interactive Communication, 2006. ROMAN 2006. The 15th IEEE International Symposium on*, pp. 39–44. IEEE (2006)
8. Harel, D., Politi, M.: *Modeling reactive systems with statecharts: the STATEMATE approach*. McGraw-Hill, Inc. (1998)
9. Hartholt, A., Traum, D., Marsella, S.C., Shapiro, A., Stratou, G., Leuski, A., Morency, L.P., Gratch, J.: All together now. In: *Intelligent Virtual Agents*, pp. 368–381. Springer (2013)
10. Hastie, H.W., Johnston, M., Ehlen, P.: Context-sensitive help for multimodal dialogue. In: *Proceedings of the 4th IEEE International Conference on Multimodal Interfaces*, p. 93. IEEE Computer Society (2002)
11. Johnston, M., Bangalore, S., Vasireddy, G., Stent, A., Ehlen, P., Walker, M., Whittaker, S., Maloor, P.: Match: An architecture for multimodal dialogue systems. In: *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, pp. 376–383. Association for Computational Linguistics (2002)
12. Jurčiček, F., Dušek, O., Plátek, O., Žilka, L.: Alex: A statistical dialogue systems framework. In: *Text, Speech and Dialogue*, pp. 587–594. Springer (2014)
13. Lamere, P., Kwok, P., Gouvea, E., Raj, B., Singh, R., Walker, W., Warmuth, M., Wolf, P.: The CMU SPHINX-4 Speech Recognition System. In: *Proc. of the ICASSP'03*. Hong Kong, China (2003)
14. Lison, P.: *Structured probabilistic modelling for dialogue management*. Ph.D. thesis, University of Oslo (2013)
15. López-Cózar, R., Callejas, Z., Griol, D., Quesada, J.F.: Review of spoken dialogue systems. *Loquens* **1**(2), e012 (2015)
16. van Meggelen, J., Smith, J., Madsen, L.: *Asterisk: The Future of Telephony*. O'Reilly, Sebastopol, USA (2009)
17. Minessale, A., Schreiber, D.: *FreeSWITCH Cookbook*. Packt Publishing Ltd (2012)
18. Neßelrath, R., Alexandersson, J.: A 3d gesture recognition system for multimodal dialog systems. In: *6th IJCAI Workshop on Knowledge and Reasoning in Practical Dialogue Systems*, pp. 46–51 (2009)
19. Pieraccini, R., Huerta, J.: Where do we go from here? research and commercial spoken dialog systems. In: *6th SIGdial Workshop on Discourse and Dialogue* (2005)
20. Povey, D., Ghoshal, A., Boulianne, G., Burget, L., Glembek, O., Goel, N., Hannemann, M., Motlicek, P., Qian, Y., Schwarz, P., Silovsky, J., Stemmer, G., Vesely, K.: The Kaldi speech recognition toolkit. In: *Proc. of the ASRU*. Hawaii, USA (2011)

21. Prylipko, D., Schnelle-Walka, D., Lord, S., Wendemuth, A.: Zanzibar OpenIVR: An Open-Source Framework for Development of Spoken Dialog Systems. In: Proc. of the TSD. Pilsen, Czech Republic (2011)
22. Ramanarayanan, V., Suendermann-Oeft, D., Ivanov, A., Evanini, K.: A distributed cloud-based dialog system for conversational application development. In: 16th Annual SIGdial Meeting on Discourse and Dialogue (SIGDIAL 2015), Prague, Czech Republic (2015)
23. Schnelle-Walka, D., Radomski, S., Mühlhäuser, M.: JVoiceXML as a Modality Component in the W3C Multimodal Architecture. *Journal on Multimodal User Interfaces* (2013)
24. Schröder, M., Trouvain, J.: The german text-to-speech synthesis system mary: A tool for research, development and teaching. *International Journal of Speech Technology* **6**(4), 365–377 (2003)
25. Suendermann-Oeft, D., Ramanarayanan, V., Teckenbrock, M., Neutatz, F., Schmidt, D.: HALEF: an open-source standard-compliant telephony-based modular spoken dialog system—A review and an outlook. In: Proc. of the IWSDS Workshop 2015, Busan, South Korea (2015)
26. Swartout, W., Artstein, R., Forbell, E., Foutz, S., Lane, H.C., Lange, B., Morie, J.F., Rizzo, A.S., Traum, D.: Virtual humans for learning. *AI Magazine* **34**(4), 13–30 (2013)
27. Swartout, W., Traum, D., Artstein, R., Noren, D., Debevec, P., Bronnenkant, K., Williams, J., Leuski, A., Narayanan, S., Piepol, D., et al.: Ada and grace: Toward realistic and engaging virtual museum guides. In: *Intelligent Virtual Agents*, pp. 286–300. Springer (2010)
28. Taylor, P., Black, A., Caley, R.: The Architecture of the Festival Speech Synthesis System. In: Proc. of the ESCA Workshop on Speech Synthesis. Jenolan Caves, Australia (1998)
29. Yu, Z., Bohus, D., Horvitz, E.: Incremental coordination: Attention-centric speech production in a physically situated conversational agent. In: 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue, p. 402 (2015)
30. Yu, Z., Ramanarayanan, V., Mundkowsky, R., Lange, P., Ivanov, A., Black, A.W., Suendermann-Oeft, D.: Multimodal HALEF: An Open-Source Modular Web-Based Multimodal Dialog Framework. In: Proc. of the IWSDS Workshop 2016, Saariselka, Finland (2016)