



HAL
open science

End-to-End Spoken Language Understanding: Performance analyses of a voice command task in a low resource setting

Thierry Desot, François Portet, Michel Vacher

► **To cite this version:**

Thierry Desot, François Portet, Michel Vacher. End-to-End Spoken Language Understanding: Performance analyses of a voice command task in a low resource setting. *Computer Speech and Language*, 2022, 75, pp.101369. 10.1016/j.csl.2022.101369 . hal-03727285

HAL Id: hal-03727285

<https://hal.science/hal-03727285>

Submitted on 22 Jul 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

End-to-End Spoken Language Understanding: Performance analyses of a voice command task in a low resource setting

Thierry Desot, François Portet¹, Michel Vacher

*Univ. Grenoble Alpes, CNRS, Grenoble INP, LIG F-38000, Grenoble, France
thierry.desot1@gmail.com, francois.portet@imag.fr, michel.vacher@imag.fr*

Abstract

Spoken Language Understanding (SLU) is a core task in most human-machine interaction systems . With the emergence of smart homes, smart phones and smart speakers, SLU has become a key technology for the industry. In a classical SLU approach, an Automatic Speech Recognition (ASR) module transcribes the speech signal into a textual representation from which a Natural Language Understanding (NLU) module extracts semantic information. Recently End-to-End SLU (E2E SLU) based on Deep Neural Networks has gained momentum since it benefits from the joint optimization of the ASR and the NLU parts, hence limiting the cascade of error effect of the pipeline architecture. However, little is known about the actual linguistic properties used by E2E models to predict concepts and intents from speech input. In this paper, we present a study identifying the signal features and other linguistic properties used by an E2E model to perform the SLU task. The study is carried out in the application domain of a smart home that has to handle non-English (here French) voice commands.

The results show that a good E2E SLU performance does not always require a perfect ASR capability. Furthermore, the results show the superior capabil-

¹Thierry Desot is now with Language and Translation Technology Team (LT³), Ghent University, Groot-Brittanniëlaan 45, 9000-B Ghent, Belgium

¹Corresponding author: francois.portet@imag.fr, Laboratoire d'Informatique de Grenoble, Bâtiment IMAG, 700 avenue Centrale, 38401 St Martin d'Hères, FRANCE

ities of the E2E model in handling background noise and syntactic variation compared to the pipeline model. Finally, a finer-grained analysis suggests that the E2E model uses the pitch information of the input signal to identify voice command concepts. The results and methodology outlined in this paper provide a springboard for further analyses of E2E models in speech processing.

Keywords: Natural Language Understanding (NLU), Spoken Language Understanding (SLU), Explainable AI (XAI), End-to-end (E2E) Learning, Smart Home

1. Introduction

Spoken Language Understanding (SLU) (Tur & De Mori, 2011) is a crucial task of most modern systems interacting with humans by speech. SLU refers to the ability of a machine to extract semantic information from speech signals in a form which is amenable for further processing including dialogue, voice commands, information retrieval, etc.

With the emergence of smart assistants in computers, smart phones and smart speakers, SLU is not only a great research area but has also become a key technology for the industry, as evidenced by several challenges including major companies such as Amazon with its Alexa prize of several million dollars to be distributed to university teams demonstrating ground breaking progress in spoken conversational AI³.

From a computing perspective, the classical approach to SLU since the early age of spoken language systems (Sears, 1988; Hemphill et al., 1990) has been a pipeline of Automatic Speech Recognition (ASR) – to transcribe the speech signal into a textual representation – which feeds a Natural Language Understanding (NLU) module – to extract semantic labels from the transcription. The main problem of such an approach is the cascading error effect which shows that any error at the ASR level has a dramatic impact on the NLU part. This is why

³<https://developer.amazon.com/alexaprize>

20 a large set of different approaches has been proposed to take into account the ASR hypotheses uncertainty within the NLU module to improve the robustness of the whole chain. This ‘classical’ approach is the default one in most industrial applications and is still an active research area (Simonnet et al., 2017).

However, since Qian et al. (2017), Neural End-to-End SLU (E2E SLU) has emerged to benefit both from the performing architecture of Deep Neural Networks (DNN) and from the joint optimisation of the ASR and NLU parts. Although the classical pipeline SLU model is still competitive, E2E SLU research has shown that the joint optimisation is an efficient way to handle the problem of cascading errors (Serdyuk et al., 2018; Desot et al., 2019a,b; Ghannay et al., 2018). In particular, it has been shown that perfect ASR transcriptions are *not* necessary to predict intents and concepts (Ghannay et al., 2018).

However, E2E SLU is more than just a way to perform joint optimisation. Indeed, in Neural End-to-End SLU, contrary to the classical pipeline, the decision stage has a direct access to the acoustic signal (e.g. prosody features). Therefore, an important question is to know **which signal characteristics (and other linguistic properties) are effectively used by the E2E SLU model**. To the best of our knowledge, there has been little research to shed light on this important question. Therefore, the goal of this study is to perform a comprehensive analysis of the linguistic features and abilities that are better exploited by E2E SLU than pipeline SLU. In particular, this paper addresses the following research questions:

1. **Can the cascade error effect of the pipeline SLU approach be avoided?** Although, this has been shown in other research, there are still too few research projects to take the answer as granted. Hence we present a comparison between pipeline and E2E SLU approach evaluated in a realistic voice command context for a language other than English: French.
2. **Is the model effectively exploiting acoustic information to perform concepts and intent prediction?** We are not aware of studies

50 having seriously explored this question. We believe that an E2E model accesses the *acoustic* levels to infer concepts and intents directly from speech. By accessing this information, the model can avoid the *cascade of errors* introduced by the interaction between the ASR and NLU models in a pipeline SLU method.

55 **3. Would an E2E SLU model be able to be more robust to variations in vocabulary and syntax?** While grammatical robustness seems to be a question only for the NLU task, we are not aware of research having investigated whether an E2E model would present a better ability to process grammatical variation. Hence, we present an acoustic and gram-
60 matical performance analysis to assess the ability of E2E SLU models to handle variation at these two levels.

Part of the comparison between pipeline and E2E SLU has been published in (Desot et al., 2019b,a). However, this paper presents updated experiments and a transfer learning approach which has not been presented yet. Furthermore, the
65 comprehensive acoustic and symbolic performance analysis was never published and constitutes the core of this article.

In this paper, Section 2 gives a brief overview of the state-of-the-art dedicated to pipeline and E2E SLU as well as the few studies deeply analysing E2E SLU performances. Our whole approach is described in Section 3 where we recap
70 the motivations, the baseline pipeline SLU and the target E2E SLU approaches before introducing the evaluation strategy. Section 4 summarises the artificial speech training data generation and describes the held out real speech test set acquired in a real smart home. Section 5 presents the results of the experiments with the baseline SLU and E2E SLU trained by transfer learning. It reports
75 superior performances for the E2E SLU model despite its lower ability at the speech transcription task than the pipeline model. Sections 6 and 7 present respectively the analysis at the acoustic level and at the grammatical level. The E2E SLU model has shown to be more robust than the pipeline one at the noise and grammatical variation levels while its ability to benefit from prosodic infor-

80 mation is less clear. These findings are discussed in Section 8 before reaching
the conclusion in Section 9.

2. Related work

Although there is a rising interest in End-to-End (E2E) SLU that jointly per-
forms ASR and NLU tasks, the E2E models have still not definitely superseded
85 pipeline approaches.

2.1. Pipeline SLU

A typical SLU pipeline approach is composed of an ASR and a NLU module.
ASR output hypotheses are fed into an NLU model aiming to extract the mean-
ing from the input transcription. The main problem with such an approach is
90 the dependence on the transcription output from the ASR module causing error
propagation and reducing the performance of the NLU module.

Hence, to deal with the uncertainty conveyed by the ASR, several methods
incorporate the handling of *N best hypotheses*. For instance, in He & Young
(2003), the ASR module (HMM) is followed by an NLU module using a Hidden
95 Vector State Model (HVS) for concept prediction. A rescoring is applied to the
N-best word hypotheses from word lattices as output from the speech recogniser.
Parse scores from the semantic parser are then combined with the language
model likelihood.

Another strategy to decrease error propagation is the use of *confidence mea-*
100 *sures*. These were used by Sudoh et al. (2006) for concept prediction, augment-
ing Japanese ASR transcriptions with concept labels. In their case, a concept
label was associated to ASR transcriptions by an SVM model only if confidence
measures were above a certain threshold. N-best list hypotheses and ASR out-
put confidence measures were also exploited using weighted voting strategies
105 (Zhai et al., 2004). Since the n^{th} hypothesis transcription contains more errors
than the $n-1^{\text{th}}$ hypothesis, voting mechanisms were used to improve perfor-
mance. For instance, a concept was considered correct if it was predicted in
more than 30% of the n -best hypotheses per reference sentence.

A third strategy is to use *word confusion networks*. For instance, Hakkani-
110 Tür et al. (2006) improve the transition between ASR and SLU concept predic-
tion, using word confusion networks obtained from ASR word lattices instead
of simply using ASR one-best hypotheses (Mangu et al., 2000). Word confusion
networks provide a compact representation of multiple aligned ASR hypothe-
ses along with word confidence scores. Their transitions are weighted by the
115 acoustic and language model probabilities. More recently, acoustic word em-
beddings for ASR error detection were trained through a convolutional neural
network (CNN) based ASR model to detect erroneous words (Simonnet et al.,
2017). This approach was combined with word confusion networks and pos-
terior probabilities as confidence measures, for concept prediction. Output of
120 the ASR model was fed to a conditional random fields (CRF) model and an
attention-based RNN NLU model. Finally, another approach, in the specific
case of smarthomes, was to use contextual information (sensors) to interpret the
ASR output using either a probabilistic network ? or a CNN-based network ?.

In Liu et al. (2020), SLU models using word confusion networks are compared
125 with 1-best hypothesis and N-best lists (N=10) for concept label and value
prediction. The ASR posterior probabilities are integrated in a pretrained BERT
based SLU model (Devlin et al., 2019). The word confusion network is fed into
the BERT encoder, and integrated into vector representation. The output layer
is a concept label and value classifier.

130 Finally, another strategy, particularly adapted when aligned labels are miss-
ing (e.g. different from an aligned BIO scheme), is to use a *sequence generation*
instead of a *sequence labelling* approach. In (Desot et al., 2019b,a), *unaligned*
NLU data was used to train a BiLSTM seq2seq attention-based model as NLU
module. Despite a lower prediction accuracy than aligned models, it provides
135 the flexibility to infer slot labels from imperfect transcriptions and speech with
disfluencies.

2.2. E2E SLU

Only recently SLU is conceived as a joint processing of the ASR and NLU tasks which decreases the error propagation between the ASR and NLU modules. Furthermore, such a model has *access to the acoustic and prosodic levels*, which can have a positive impact on the performance of SLU. For instance, Serdyuk et al. (2018) trained a sequence-to-sequence (seq2seq) model on clean and noisy speech data to infer intents directly from audio MFCC. Such an approach showed that some prosodic aspects of the speech signal were exploited by the E2E model for intent classification (e.g. question vs imperative voice).

E2E SLU is also driven by the intuition that recognising speech, *word by word*, is *not* necessary. In (Ghannay et al., 2018), the Baidu Deep Speech ASR system (Hannun et al., 2014) was trained on transcriptions enriched with concept labels. Eight concepts were injected into the ASR transcriptions as symbolic labels. In order to reduce the importance that the connectionist temporal classification (CTC) cost function (Graves et al., 2006; Watanabe et al., 2017; Ueno et al., 2018) assigns to each character and to draw more attention to the concept symbols, all character sequences not related to a concept label were replaced by one and the same symbol.

Different from these E2E models that predict intent and concept labels directly from speech, a *transfer learning* technique allows the training of a complete E2E model through sub-tasks (for instance, forcing hidden layers to predict phonemes), thereby providing an easier learning path. Combined with a curriculum learning that presents the easy examples before the more complex ones during training, convergence of the learning algorithms is accelerated (Krueger & Dayan, 2009). For instance, in Lugosch et al. (2019), a transfer learning for intent prediction is applied by training first an ASR model and then adapting it to an SLU task to predict concepts that are finally mapped to intents.

In Caubrière et al. (2019, 2020), this type of E2E SLU is performed for concept prediction, using the Baidu Deep Speech ASR tool. A phase of training an ASR model, is followed by three phases of learning concepts with an increasing complexity. The approach showed a clear gain in performance compared with a

classical pipeline approach on the French Media corpus.

2.3. E2E SLU analysis

170 Although E2E SLU is emerging as an alternative approach for pipeline SLU
systems, we were not able to find research offering an in-depth performance
analysis. In Caubrière et al. (2019), an *Error analysis* of their E2E SLU system
is performed. They show that concept deletion errors are not mainly caused by
the ASR capability of the system, but occur as a consequence of a segmentation
175 problem. On top of that, unseen concepts are better predicted using a transfer
learning approach. In the study of Rao et al. (2020), a pipeline SLU system
is compared with an E2E SLU model for Amazon Alexa, where the interface
between ASR and NLU is a shared 1-best hidden layer. They show that joint
ASR and NLU training improves SLU performances for ASR erroneous output
180 transcriptions that impact NLU performances in a pipeline model.

In Denisov & Vu (2020), a pretrained ESPnet ASR model encoder was com-
bined with transformer based pretrained contextual BERT embedding. They
analysed the fine-tuning of the E2E SLU *layers*. Results indicate that fine-
tuning the ASR encoder layers is more beneficial than the NLU layers. This
185 would mean that the acoustic representation must be adapted to the concept
extraction task and is thus different from the ASR task.

This brief state of the art shows that handling the cascading error effect
has been the main focus of pipeline SLU studies and the main motivation for
E2E systems. Despite a gain of performance of E2E SLU and some studies
190 analysing the potential effect of such gain, the impact of the E2E SLU model's
access to the acoustic level has not been investigated in-depth. To the best of
our knowledge, we are not aware of any other study analysing the impact of
acoustic features on E2E SLU performances and the robustness of such a model
to grammatical mismatch.

195 **3. Method**

Before presenting the overall approach of the paper, we recall the research questions mentioned in the introduction.

200 The state-of-the-art section showed that one of the main drawbacks of a pipeline SLU is the cascade of error effect. Hence, the main objective of previous and recent SLU approaches was to reduce the impact of ASR errors on NLU performances. Our goal in this study is not only to avoid this cascade of errors but also to **understand what advantages an end-to-end (E2E) SLU approach can offer over a traditional pipeline approach** (this will be detailed in Section 5).

205 Given that the E2E approach extracts semantics directly from the acoustic signal, an interesting research question we address in this paper is **whether the E2E model exploits prosodic information to infer intent and concepts from the speech signal** (this will be the subject of Section 6).

210 Another drawback of the classical ASR and NLU systems that can impact the NLU part negatively is the problem of out of vocabulary words (OOV) and unusual syntactic structures. By modelling linguistic phenomena at a finer granularity, we want to verify **whether an E2E SLU model is more robust to vocabulary and syntactic variations** (section 7 will detail this study).

215 This section details the overall approach we have followed to address the above challenges.

3.1. Overall approach

The solution we propose consists in considering the SLU problem as a slot-filling problem applied to the field of voice command in a smart home.

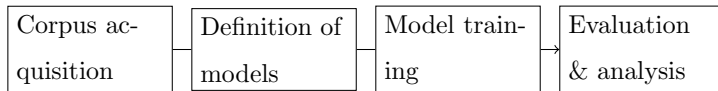


Figure 1: Overview of SLU development steps

Figure 1 describes the steps of the approach. The first step was to collect a
220 *test corpus* to evaluate the SLU approaches on realistic data. To solve the lack
of *training* data, we have chosen an expert based artificial corpus generation
approach. In order to reproduce a realistic situation, no data from the test
corpus was used to train the models.

Then, we defined two baseline systems: the pipeline SLU and E2E SLU. *The*
225 *pipeline SLU* is a combination of 2 state-of-the-art ASR and NLU modules. *The*
End-to-end SLU model is a pyramidal RNN multi-task model that combines a
CTC cost function and an attention based encoder and decoder . Theoretically,
this type of model is capable of handling OOV words, it is therefore interesting
to study whether the interaction between attention and CTC can strengthen
230 the robustness of such a model on test data with high *linguistic variability*. The
SLU objective and the two SLU baselines are further introduced in the following
section 3.2.

Once the models defined, they were trained on the generated corpus. The
pipeline training is detailed in Section 5.1 and the E2E training in Section 5.2.

235 Finally, once the models have been trained for the task, we performed various
evaluations to assess to which extent the results are correlated to external factors
such as noise condition, gender, pitch variation, syntactic complexity, etc. The
correlation measures used for the study are introduced in Section 3.5 while the
experiments are detailed in Sections 6 and 7.

240 3.2. SLU approach

The target of our SLU experiments are commands without linguistic context
and with one intent per utterance. The notion of intent is close to that of a
speech act which is the speaker’s communicative activity in which an utterance
produces an effect on its interlocutor (Crystal, 2011). In this study, the intent
245 is the type of act addressed to the home automation system. For example, the
statement “Allume la lumière” (*turn on the light*) conveys the intent to the smart
home to change the state of an object. To characterise the voice command, it is
also necessary to identify its *concepts* or *slots* representing the most important

information (i.e. entities and actions). This process is called *slot-filling* (Tur & De Mori, 2011). Figure 2 shows the two ways of extracting intent and slots we followed in this paper. In the pipeline case (on the top of the figure), the input utterance is first transcribed by the ASR module to be analysed by an NLU module that generates a sequence of **concepts** that are supposed to be found in the speech input. In the example, the intent is **set_device** while the concepts are the action with value **turn_on** and the device with value **light**. In the E2E case (on the bottom of the figure), the SLU target is seen as an enriched transcription task. The SLU model is trained to surround the transcribed words with specific characters such as **^** to delimit an action or **}** to delimit a device. Intents are classified using a special character at the beginning and end of a transcription (here **@** means the **set_device** intent).

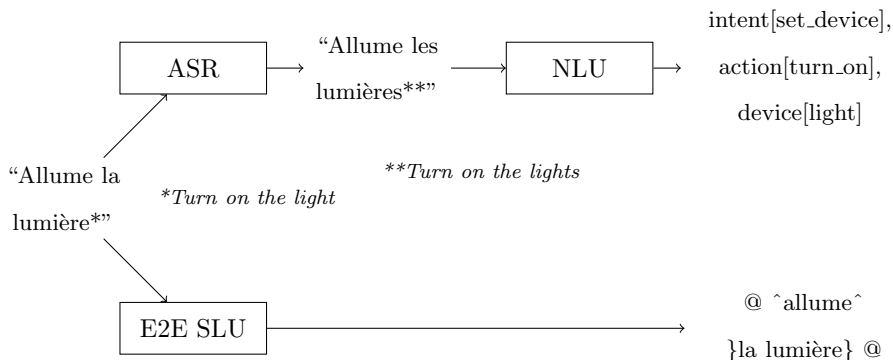


Figure 2: Comparison of pipeline and E2E SLU tasks

3.3. Baseline pipeline SLU

As in (Desot et al., 2019b), the ASR component of our pipeline SLU is the Kaldi tool, `mnet2` version. This neural-network ASR training framework allows training with large amounts of data using multiple GPUs or multi-core machines. It uses speaker adapted features from the GMM (Gaussian Mixture Model) system, so a first pass of GMM decoding and adaptation are required (Povey et al., 2011, 2015).

Mel Frequency Cepstral Coefficients (MFCC) were used as input features. Kaldi also allows using several adaptation methods of the acoustic models to the speaker, such as *Maximum Likelihood Linear Regression* (MLLR) (Leggetter & Woodland, 1995), *Constrained Maximum Likelihood Linear Regression* (fMLLR) (Digalakis & Neumeyer, 1996) and *Speaker Adaptive Training* (SAT) (Anastasakos et al., 1996). As the ASR component has to interact with the NLU module in a pipeline system in the real time setting of a smart home, the mnet2 *online* version was also used.

Regarding the NLU module, we approach it as a *sequence generation task* with *unaligned* data. The SLU task is seen as a translation problem where the input must be abstracted to generate output intent classes and concept labels. For that reason, the NLU module was a seq2seq bidirectional LSTM encoder and decoder attention-based model. This was our strategy to decrease errors of our baseline pipeline SLU model, due to the imperfect transcription output of the ASR component that impacts the NLU. Using an unaligned approach the model should learn to associate several words to one slot label without aligned data. Furthermore, classical BIO alignment cannot be assumed for pipeline and E2E SLU when input data consists of spontaneous speech with disfluencies that often cause ASR deletion and insertion errors. Hence, a robust NLU subpart is needed that can handle those and that can be trained with unaligned labels. In Mishakova et al. (2019), we showed that such an NLU approach is competitive with state of the art *aligned* NLU CRF models (Jeong & Lee, 2008) and also with DNN-based models (Mesnil et al., 2015; Bapna et al., 2017; Liu & Lane, 2016; Huang et al., 2017) that deal with the NLU problem as a *sequence labelling task*.

3.4. End-to-end SLU

The E2E approach as outlined in Desot et al. (2019a) was based on the ESPnet ASR toolkit (Watanabe et al., 2018). It integrates the Kaldi data preparation, extracts Mel filter-bank features, and combines Chainer and PyTorch deep learning tools (Tokui et al., 2015; Paszke et al., 2017). The encoder consists of

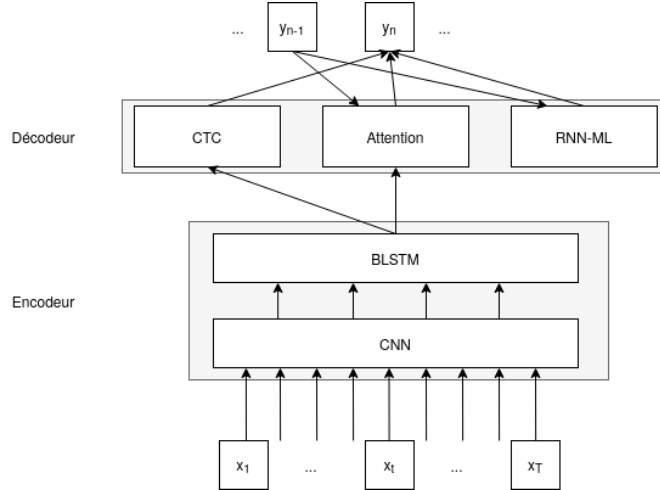


Figure 3: ESPnet architecture with the different training/inference strategies (CTC, attention, ML RNN)

a very deep convolutional neural network (VGG) followed by six bidirectional pyramidal subsampling bi-LSTM layers. Figure 3 includes an overview of the
 300 ESPnet architecture.

Convolutional neural networks (CNNs) have achieved great success in image recognition (Cho et al., 2018). In the context of ASR, CNNs are usually used as feature extractors, while the HMM part is typically replaced by RNNs that provide a distribution over sequences directly (Zhang et al., 2016). The success
 305 of using CNNs in ASR tasks, can be attributed to the use of local filtering and maxpooling in the CNN architecture. This combination turns out to be a better strategy than a GMM model that represents the entire frequency spectrum as a whole. Another benefit is better robustness against ambient noise. In order to locally apply filtering, a frequency scale is needed that can be divided into
 310 a number of local bands. Therefore, MFCC features are not suitable because of DCT-based (Discrete Cosine Transform) decorrelation transform. Indeed the Gaussian mixture model needed decorrelated input feature dimensions while CNNs can benefit from correlated inputs. For that reason, filter-bank features are more fit for local filtering using CNNs. Max pooling and filtering is also

315 an alternative for SAT (Speaker Adaptive Training) and MLLR (Maximum Likelihood Linear Regression) that transform speech features into a canonical speaker space (Abdel-Hamid et al., 2012).

In ESPNet, mapping from acoustic features to character sequences is performed by a *hybrid* multitask learning that combines CTC (Amodei et al., 2016; 320 Graves et al., 2006) and attention (Bahdanau et al., 2014). The attention mechanism allows a more flexible alignment, which focuses on the important features and character sequences whereas the ASR alignment is monotonic. A trade-off hybrid CTC and attention-based approach finds a balance between attention and CTC.

325 For moving the ESPNet model from a transcription objective to a SLU objective, as described in Figure 2 the transcriptions were enriched with symbols representing intent classes and concept labels. This strategy has been previously used with success by Ghannay et al. (2018); Desot et al. (2019a,b).

3.5. Analysis measures

330 An E2E SLU approach has access to the *acoustic* and *prosodic* information of the input signal. Therefore, in the analysis we intend to assess to which extent the model is capable of exploiting paralinguistic information to infer semantic information. Table 1 gives an overview of the analysis levels that we consider in the study.

Table 1: Analysis levels of ASR, pipeline and E2E SLU on the VocADom@A4H test set

Analysis level	Acoustic			Symbolic	
	Source		Prosody	Lexical	Syntax
Analysis:	Noise	Gender	Avg. F0	OOV	Variation
Hypothesis:	(1) Does the E2E model benefit from prosodic, acoustic information ?				
	(2) Is the E2E model more robust to lexical and syntactic variations?				

335 The acoustic analysis was devoted to two main features : the robustness to the variability of the source (here, the background noise and gender) and the capability of the model to exploit acoustic features. The robustness to

OOV and syntactic variability was studied under the term 'symbolic' in order to distinguish it from the purely acoustic part of the study. For each study, the analysis consisted in generating or extracting specific stimuli to feed models, then measuring the *performance* of each model at various levels (intent, slot filling, speech recognition) and assessing their *correlations* with the acoustic or symbolic features of the input stimuli. We describe below which performance and correlation measures we used in the study.

For assessing the intent prediction performance, since the task consists in choosing for each utterance one possible intent among a restricted set of classes, we therefore used classic measures in classification, recall = $\frac{TP}{TP+FN}$ and precision = $\frac{TP}{TP+FP}$, where TP is *True Positive*, FP is *False Positive*, and FN *False Negative*. *Precision* expresses the proportion of correctly predicted intents in the set of *predictions*, whereas *recall* expresses the rate of correct predictions among the set of instances to be predicted. The F-measure or F1 score provides a single score that balances both precision and recall and is calculated as follows,

$$F1 = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (1)$$

For slot-filling performance, we used the concept error rate (CER) which is defined as the ratio of the sum of deleted, inserted and confused concepts w.r.t. a Levenshtein alignment for a given reference concept string (Hahn et al., 2008). In this paper, we calculated the CER in a similar way, but we did not take the label sequence order into account since we used a generative approach. Hence, a reference sequence **action**, **device** provides the same information as a hypothesis such as **device**, **action**. As for the E2E approach, transcriptions are enriched with symbols that represent concepts and intents. We only consider the symbol sequences of reference and hypothesis concepts. Figure 4 shows how the labels are extracted from the outputs to compute the CER. In this example, the hypothesis transcription shows an erroneous deletion of the symbol } and obtains thus a CER score of 50%.

For the ASR level, performances were evaluated using the WER. The WER

Reference = }la lumière} doit être ^allumée^ => }^
(the light must be turned on)

Hypothesis = ^allume^ le soir => ^
(turn on in the evening)

CER = 50%

Figure 4: E2E SLU - calculation of the concept error rate that is symbolically represented

is defined as the sum of *insertions* (I), *deletions* (D) and *substitutions* (S) of words compared to the number of words N in a reference transcription that has manually been verified. The alignment between the reference and hypothesis transcription is obtained by dynamic programming in order to find the alignment leading to the minimal WER. The WER is calculated as follows :

$$\text{WER} = \frac{I + S + D}{N} \times 100 \quad (2)$$

As a final evaluation step, our main strategy is to compute the correlation of the resulting SLU and ASR performances with respect to the characteristics of the input stimuli. We measure *Pearson* and *Spearman's rank* correlations between performances (CER, WER, F1) and input features (e.g. *energy*).

Pearson's correlation coefficient r measures the strength of the association between two variables x and y assuming a normal distribution of values. Pearson's correlation is calculated as follows, where n denotes the number of elements :

$$r = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sum_{i=1}^n \sqrt{(x_i - \bar{x})^2} \sum_{i=1}^n \sqrt{(y_i - \bar{y})^2}} \quad (3)$$

r varies between -1 and 1. There is no correlation when r is equal to 0. 1 indicates a strong correlation and -1 a strong negative correlation.

Spearman's rank correlation, r_s , measures correlations between *ranked variables* and is calculated as follows where n denotes the number of elements :

$$r_s = 1 - \frac{6 \sum_{i=1}^n (x_i - y_i)^2}{n(n^2 - 1)} \quad (4)$$

In order to determine if the resulting correlation coefficient is significant, the p -value is often used. It is calculated in hypothesis testing to determine whether or not to reject a null hypothesis. The p -value for the Pearson or Spearman correlation coefficient ($coef$) uses the distribution law t ,

$$t = r \sqrt{\frac{n-2}{1-coef^2}} \quad (5)$$

We reject the null hypothesis, $H_0 : coef = 0$, if the p -value is less than 0.05 ($p < 0.05$). The p -value or hypothesis $H_1 : coef \neq 0$ (two tails test), is computed as follows :

$$p = 2 * P(T > |t|), \quad (6)$$

where P denotes the probability and T follows a distribution t with $n - 2$ degrees of freedom. We distinguish between correlations for which $p < 0.05$ that we mark with * and correlations for which $p < 0.01$ that we denote with **.

In order to measure and compare the robustness of the pipeline and E2E SLU models to lexical and grammatical variations (see Section 7), we verified the impact of an increased OOV (Out Of Vocabulary words) on SLU performances by gradually replacing domain specific vocabulary by synonyms that do *not* occur in the training set. We also measured the impact of the syntactic variability in the speech of our target users by predicting concepts and intents of the test data where we inserted syntactic structures and disfluencies that hardly occurred in the training set (Table 1). Our hypothesis is that the learning of the E2E SLU model, that combines the CTC and attention mechanisms can enhance the robustness of the E2E model to *linguistic variability*.

4. Data

To acquire target corpora in sufficient size to train deep neural network models, we used a combination of realistic data, synthetic data and out-of-domain data.

Realistic data was acquired within a real smart home with naive users, using a Wizard-of-Oz strategy to acquire diverse and contextualised voice commands. Indeed, since our target users are senior adults, they tend to deviate from too strict a grammar (Möller et al., 2008; Takahashi et al., 2003; Vacher et al., 2015),
410 hence the need for a realistic corpus that accounts for the rich set of possible sentences and pronunciations. This corpus is quickly introduced in section 4.1 and was made available to the community (Portet et al., 2019).

Acquiring such kind of realistic corpus is highly time consuming and leads to
415 an amount of data far too small for machine learning. We tackled this problem by automatically generating a domain-specific *synthetic speech training* corpus using Natural Language Generation controlled by the semantic space of the smart home. This synthetic generation is described Section 4.2.

Finally, since the artificial generation presents a good semantic coverage but
420 a poor diversity, we also collected other out-of-domain corpora that can be used either to enrich the training data or to perform transfer learning. This collection of corpora is listed in Section 4.3.

4.1. VocADom@A4H test data

The VocADom@A4H corpus (Portet et al., 2019; Desot et al., 2018) includes
425 about twelve hours of speech data and was acquired in realistic conditions in the Amiqua4Home smart home⁴. Eleven participants uttered voice commands while performing activities of daily living for about one hour in different rooms including a kitchen, a living room, a bedroom and a bathroom. Out-of-sight experimenters reacted to participants' voice commands following a wizard-of-Oz
430 strategy to add naturalness to the corpus. Furthermore, experimenters were also present in the home to act as visitors of the participant. At the end, the corpus consists of a mixture of spontaneous and read voice commands, with lexical and syntactic variation. Some of the utterances were recorded with background noise (use of vacuum cleaner, radio, tv etc.).

⁴<https://amiqual4home.inria.fr>

Table 2: VocADom@A4H test data overview

VocADom@A4H	utterances	words	intents	slot labels
intents(1)	2612	430	7	14
none intents(2)	4135	1326	1	-
complete(3)	6747	1462	8	14

435 Each voice command was prefixed with a keyword (chosen by the participant
in a list) to activate the smart home (e.g. “Minouche, lower the blind of the
bedroom” where ‘Minouche’ was one of the possible keywords). The partici-
pants’ and experimenters’ speech was semi-automatically transcribed and then
corrected manually. For SLU, data was manually annotated with intent classes
440 and slot labels whose semantics was defined in accordance with the smart home
capabilities. A description of the semantic labels for the slot is provided in
Figure 5.

At the end of the annotation process, 6,747 utterances constitute the dataset.
As shown in Table 2, it consists of voice commands (*intents(1)*), and utterances
445 other than voice commands (*none intents(2)*). This *realistic* corpus is the held
out test set used for all our SLU experiments. It is freely available for research
purpose at vocadom.imag.fr. For more information about this corpus, the reader
is referred to (Portet et al., 2019).

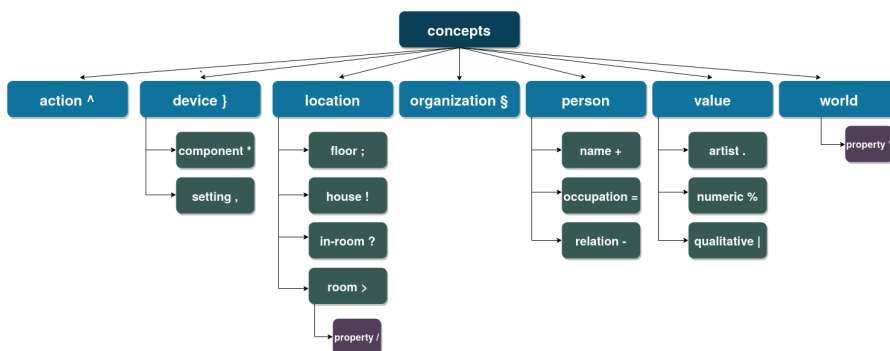


Figure 5: Tree structure of concepts with symbolic representation

4.2. Automatic artificial training data generation

450 To gather a large amount of data with a broad coverage, we used an expert-based Natural Language Generation (NLG) approach (Gatt & Krahmer, 2018). An NLG approach is more easily controlled compared to a constrained RNN language model for data augmentation (Hou et al., 2018). The generation was performed in two phases:

- 455 1. Voice command utterances were generated as text and semantically annotated at the same time.
2. Generated textual voice commands were fed to a speech synthesizer to provide a complete corpus of speech annotated with semantic information.

The NLG system was based on the NLTK python library feature-based context-free grammar (FCFG) (Bird et al., 2009) allowing for sentence generation, and for features (i.e. slot information) to be attached to the final output sentences. The grammar defines intents as a composition of their possible constituents, with constraints on generation.

The semantic space consisted of four general intent classes:

- 465 • **contact** which allows a user to place a call;
- **set** to make changes to the state of objects in the smart home;
- **get** to query the state of objects as well as properties of the world at large;
- and **check** to check the state of an object. Slot labels are divided into eight categories: the **action** to perform, the **device** to act on, the **location** of the device or action, the **person** or **organization** to be contacted, 470 a device **component**, a device **setting** and the **property** of a location, device, or world.

Syntactical variation was also part of the grammar design. Similar to the test set, each voice command includes a keyword to activate the smart home. 475 Maximising all combinations of semantic labels that result in meaningful utterances, the grammar generated about 77k phrases. These generated sentences

were automatically annotated with 17 different concepts and eight different intent classes, based on the general categories defined by our semantic space. Table 3 includes an overview of intents. An overview of concepts is presented in Figure 5, while Table 3 provides a comparison of the number of utterances per intent between the artificial and realistic corpus. The generation process and its evaluation is detailed in (Desot et al., 2018).

Table 3: Distribution of utterances broken down by intent for the Artificial (Artif.) train and VocADom@A4H (Real.) test corpus with examples

Intent	Symb.	Example utterances	#	
			Artif.	Real.
Check_device	#	<i>minouche is the window open ?</i>	2754	284
Contact	[<i>vocadom call a doctor</i>	567	114
Get_room_property	{	<i>berenio what's the temperature ?</i>	9	3
Get_world_property]	<i>ulyse what's the time ?</i>	9	3
None		<i>the window is open</i>	-	4135
Set_device	@	<i>hestia lower the blinds</i>	63,288	2178
Set_device_property	-	<i>ichefix decrease the TV volume</i>	7290	9
Set_room_property	&	<i>chanticou decrease the temperature</i>	3564	21

The semantic annotation part of the synthetic corpus was generated in two versions: one for the pipeline and one for the E2E approach. Table 4 provides examples of these two formats. For the E2E SLU approach, the artificial corpus transcriptions are enriched with intent class and slot label symbols (Desot et al., 2019b,a). A similar approach was applied in (Ghannay et al., 2018), however our transcriptions are enriched with *both* intent and slot label symbols.

As an SLU system extracts slot labels and intent classes from speech, we

Table 4: Artificial corpus format for the generative approach (pipeline) and the enriched transcriptions approach (E2E)

<p>Format for generative NLU</p> <p>Pipeline SLU (“vocadom close the door”)</p> <p>(Source) vocadom ferme la porte</p> <p>(Target) intent[set_device], action[close], device[door]</p>
<p>Format for symbolically enriched transcription</p> <p>E2E SLU (“vocadom switch on the light”)</p> <p>(Source + Target labels injected)</p> <p>@ VocADom ^allume^ }la lumière} @</p> <p>SET_DEVICE intent class symbol @/ Action slot symbol ^ /</p> <p>Device slot symbol }</p>

490 used a speech synthesizer to generate spoken utterances for the 77k artificial sentences, using the open-source Ubuntu SVOX ⁵ female French voice ⁶.

4.3. Collection of realistic datasets

The artificial corpus, though of great semantic coverage, does not cover the diversity of speech that can be found in the test data. For instance, Table 2
495 reports that **none** is the majority intent class in the VocADom@A4H test set. Furthermore, the artificial corpus contains only artificial speech produced by one synthetic voice. In order to increase the number of **none** intents in the artificial training data as well as to add voice diversity, the ESLO2 corpus (Serpellet et al., 2007) of conversational French speech was added to the training set. Sentences unrelated to voice command intent were extracted (i.e. **none** intent) and
500 manually filtered. Only out of domain utterances were kept for collecting **none** intent training data. Furthermore, similarly to the VocADom@A4H corpus, it contains frequent disfluencies. The small domain specific SWEET-HOME cor-

⁵<https://launchpad.net/ubuntu/+source/svox>

⁶<https://doc.ubuntu-fr.org/svoxpico>

Table 5: Comparison of SLU training and test data (OOV = test set words not seen in training data)

training set	utterances	words	intents	slots	perpl.	OOV	speech (hours)
Artif.	77,481	187	7	17	124.41	307	81.25
Sweet-Home	1412	480	6	7	49.33	343	2.5
Eslo2	161,699	29,149	1	-	151.90	211	126
Total	240,592	30,821	8	17	372.06	235	209.75

pus, with distant voice commands (?Vacher et al., 2014) was also added to the
 505 training data.

Table 5 includes an overview of the complete SLU training set consisting of the artificial corpus and the SWEET-HOME and ESLO2 corpora. The perplexity (*perpl.*) and OOV with respect to the test set are provided using a 3-gram language model learned on each corpus. It can be seen that the vocabulary of
 510 the artificial corpus is quite poor despite a good semantic coverage. Because of this small vocabulary and the strict syntactic pattern the perplexity and OOV stay high. By contrast, the SWEET-HOME corpus has a relatively low perplexity even with such a small amount of data. Finally, the large vocabulary of ESLO implies the smallest OOV rate.

515 **5. Training of the pipeline and E2E SLU models: impact of the cascade error effect**

5.1. Baseline pipeline SLU

The ASR and NLU modules of the pipeline SLU were trained separately. For the ASR module, a large acoustic model was trained using 472.65 hours
 520 of *Real Speech* data from the French corpora ESTER1 (Galliano et al., 2005) and ESTER2 (Galliano et al., 2009), REPERE (Giraudel et al., 2012), ETAPE (Gravier et al., 2012), BREF120 (Tan & Besacier, 2006), AD (Vacher et al., 2008), SWEET-HOME (Vacher et al., 2014), CIRDO (Vacher et al., 2016) and

Table 6: French corpora used to train an acoustic model using the ASR tool Kaldi

Corpus	# hours of speech
ESTER1	100
ESTER2	100
REPERE	60
ETAPE	30
BREF120	51.50
AD	0.5
SWEET-HOME	2.5
CIRDO	2
ESLO2	126
Total	472.65

the corpus of spontaneous speech ESLO2 (Serpellet et al., 2007), and also the
 525 small domain specific SWEET-HOME corpus (Table 6).

For Kaldi NNET2 (Section 3.3), an architecture of 4 hidden layers with 1024
 hidden units, a Softmax layer with 4748 units and SGD (Stochastic Gradient
 Descent) was used. The learning rate started at 0.01 and ended at 0.0001, with a
 batch size of 128. The total training lasted for 253 hours, using 1 GPU GeForce
 530 GTX TITAN Black. Acoustic features were 13-dimensional MFCC features.
 For further detail, the reader is referred to (Desot et al., 2019a,b, 2020).

For the seq2seq NLU module described in Section 5.1, input words were first
 passed to a 300-unit embedding layer. The encoder and decoder were each a
 single layer of 500 units. Adam optimizer was used with a batch size of 10,
 535 using gradient clipping at a norm of 2.0. Dropout was set to 0.2 and training
 continued for 10,000 steps with a learning rate of 0.0001. Input sequence length
 was set to 50 and output sequence length to 20, with a beam search of size 4.
 The training data was the combined semantically annotated datasets: artificial,
 SWEET-HOME and the filtered ESLO2 utterances without intent (Table 5).

540 Once ASR and NLU models trained, inference with the pipeline approach

just consisted in feeding the best ASR transcription generated using Kaldi NNET2 to the seq2seq NLU module. Table 7 shows the differences in performance between the pipeline SLU approach (Pipeline SLU) and the NLU model that are larger for slot predictions as compared to intent predictions.

Table 7: Pipeline and E2E SLU performances, % F1-score - Concept Error Rate - WER on VocADom@A4H. † with ESPNet as ASR

Model	Hours of speech	(%) TTS in train	Intent F1-score	Slot CER	WER
Pipeline:					
ASR	472.65	0.00	-	-	22.92
NLU	-	-	85.51	33.78	-
SLU	472.65	0.00	84.21	36.24	-
E2E:					
ASR	553.90	14.67	-	-	46.50
SLU	553.90	14.67	47.31	51.87	-

545 5.2. E2E SLU with transfer learning

For the E2E experiments, we used ESPnet default settings (Desot et al., 2019a) in order to train on speech data, with slots and intents symbolically injected in the transcriptions. Since the training was end-to-end, the training data was composed of the ASR training data (472.65 hrs) plus the NLU training data (81,25) were the artificial dataset was synthesised (cf. Section 4). This represents 553.9 hours of training data.

Table 7 reports the results on the VocADom@A4H test set (*E2E*, *ASR*). For the Pipeline and E2E ASR task, a far worse WER is exhibited for the E2E model. Similarly, for the SLU task, the E2E model exhibits far worse performances for CER and F1.

The large amount of data for the E2E learning, far from allowing generalisation, mainly led to unbalanced learning. In the training data, the open domain was too large to allow domain data to properly drive the training. Another way

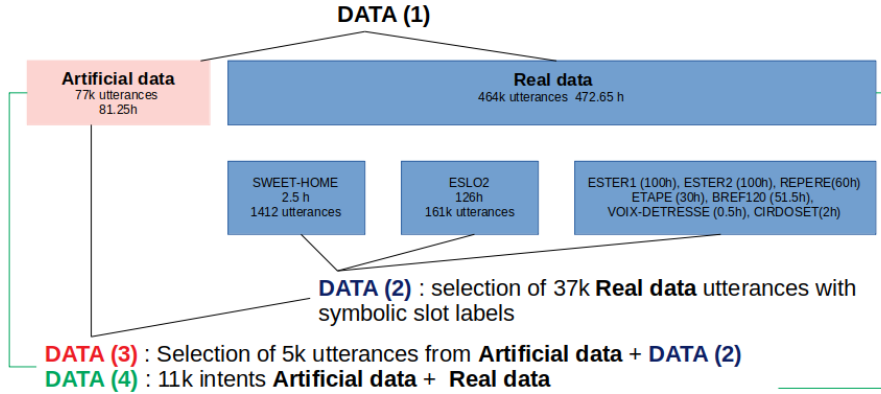


Figure 6: Transfer learning - Concept and intent prediction

to take advantage of a large non-domain-specific dataset, and a small domain
 560 specific dataset, is to use a transfer learning approach (cf. state of the art in
 Section 2). This means pre-training a large amount of speech data on an ASR
 task in order to make the model learn the input representations of the acoustic
 signal. Once the model is learned, training is restarted on an SLU task with
 other datasets designed specifically to learn concepts and intents. The complete
 565 process consists of 4 steps (Figure 6 and Table 8): 3 steps for the prediction of
 concepts and a fourth step for intent prediction:

1. For the first step an ASR model is trained (16 epochs) on the set of real
 and artificial speech utterances (553.9h) ($data(1)$);
2. For the second step, there is a training (12 epochs) on 37k real sentences of
 570 speech data, as part of $data(1)$, which contains symbolic concepts specific
 to the home automation domain ($data(2)$);
3. The third stage is a training (9 epochs) on 3800 artificial corpus sen-
 tences whose concepts are *missing* or *underrepresented* in $data(2)$.
 1651 statements were also taken from $data(2)$ containing underrepresented
 575 concepts. Indeed we combined a *transfer* learning and a technique of *data*
duplication resulting in 5451 statements ($data(3)$);
4. The final step is a training (11 epochs) on intent utterances $data(4)$ which

contains 11K statements of real and artificial speech.

Table 8: E2E SLU - intent and concept symbols

Concept (data(3))
<pre> hestia s'il vous plaît ^baisser^ }la lampe} >de la chambre> action device location-room (hestia please decrease light in the room) </pre>
Intent + Concept (data(4))
<pre> @@ hestia s'il vous plaît ^baisser^ }la lampe} >de la chambre> @@ set_device action device location-room </pre>
Intent + Concept - without words outside slots (data(4*))
<pre> @@ hestia * * * ^baisser^ }la lampe} >de la chambre> @@ set_device action device location-room </pre>

Results are presented in Table 9 which shows performance results for all
580 phases of transfer learning. The baseline results with Kaldi (*Pipeline SLU*) and
the E2E approach are shown for comparison (*E2E SLU-small* is not used as a
baseline since it includes 1k from the test set). The results for the transfer learn-
ing from the ASR (*Data(1)*) to the SLU task (*Data(2)*), (*Data(1)* \rightarrow *Data(2)*),
show lower performances than the approach with a reduced data set. In order
585 to verify that our results are *truly* based on a transfer learning, we compared
performances based on a model trained *only* on the 5k statements of *Data(3)*.
Training only on *Data(3)* shows its limits (CER amount to 69.11%). On the
other hand, a training on the SLU task *Data(2)* that is transferred to SLU
task *Data(3)* (*Data(2)* \rightarrow *Data(3)*) shows its efficiency for concept prediction,
590 performing better (CER = 32.12 %) than the baseline pipeline SLU approach
(*Pipeline SLU*) obtaining a CER of 36.24%. These results indicate the relevance
of transfer learning for E2E SLU.

For intent prediction, we continued the transfer learning principle using in-
tent data (*Data(4)*). For this intent learning task, on top of the transcriptions
595 that are augmented with concept symbols, intent symbols were inserted (Table

Table 9: E2E SLU - intent and concept symbols

Model	Intent (%) F1-score	Concept (%) CER
Training without transfer learning :		
Pipeline SLU	84.21	<i>36.24</i>
E2E SLU	47.31	51.87
E2E SLU Data(3)	-	69.11
Transfer learning E2E SLU :		
Data(1) → Data(2)	-	42.19
Data(2) → Data(3)	-	32.12
Data(2) → Data(3) → Data(4)	68.13	-
Data(2) → Data(3) → Data(4*)	<i>74.57</i>	-

8). Table 9 shows that we could not outperform the sequential SLU intent prediction performances. Best transfer learning results for intent prediction were obtained using the $Data(2) \rightarrow Data(3) \rightarrow Data(4^*)$ model where word tokens “outside concepts” have been replaced by asterisk symbols. Nevertheless the latter model outperforms the reduced model’s intent prediction (*ESpnet-small*).

5.3. ASR impact on E2E SLU

In Section 5.2, we have shown that the E2E SLU model outperforms the pipeline SLU approach, by applying transfer learning, for concept prediction in spite of ASR hypothesis transcriptions that are far from perfect. Hence, we expect a weak correlation between WER values from ASR performance and CER values from E2E SLU concept prediction, based on transfer learning. In order to verify this, we calculated the Pearson and Spearman correlation coefficients between the WER value per utterance and the CER value per utterance. Table 10 shows very significant correlations but which are not as strong as one would expect. Thus an improvement for an ASR task should have a positive impact on the SLU task but it is by far not sufficiently predictive of E2E SLU performances. This answers our first question by confirming that E2E approaches are indeed

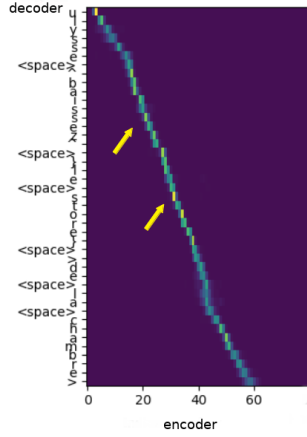


Figure 7: Utterance with attention for concept labels

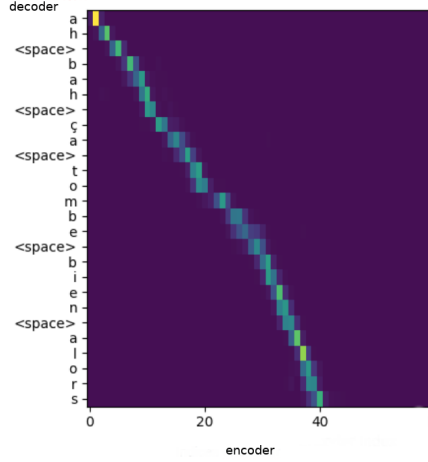


Figure 8: Utterance without concept labels

a way to diminish the cascade of error effect of the ASR task on the NLU task.

Table 10: Pearson and Spearman correlations E2E ASR - E2E SLU

Model	WER (%)	CER (%)
ASR E2E	46.50	-
E2E SLU	-	32.12
Correlation coef.		
Pearson (r)	0.26**	
Spearman (r _s)	0.25**	

* means $p < 0.05$; ** means $p < 0.01$

6. Acoustic Analysis of E2E SLU predictions

615 Unlike a sequential SLU approach, an E2E SLU approach has access to *acoustic* information of the input signal. It is therefore relevant to verify whether the model exploits paralinguistic indices to infer semantic information.

6.1. Acoustic information impact in the E2E model

One way to qualitatively analyse E2E neural model is to analyse the attention map. Using ESPnet, an attention heat map was generated for the test set utterance “ulyse baissez le store de la chambre” (*ulyse lower the blind in the room*), with concept labels **action** ‘lower’ and **device** ‘the blind’. The yellow arrows in Figure 7, pointing to the lighter color areas, show increased attention for the concept labels, especially around the hat and brace symbols that represent the concept labels **action** and **device** respectively. On top of that, pitch and energy were measured for the same utterance using *Praat*⁷. Figure 9 shows a pitch contours (blue line) that increases for the concepts.

In order to exclude that the increased attention around the symbols is caused by white-spaces (<space>), another attention heat map for a test set utterance *without* concepts was generated for the utterance “ah bah ça tombe bien alors” (*ah well that’s good then*) and clearly does not show any increased attention around white-spaces (Figure 8). This indicates that the E2E SLU model seems to learn that the concept symbols are more important than the other character symbols. This result led us to the research question whether E2E SLU benefits from *acoustic* information in order to predict concepts and intents and if correlations exist between prosody, on the one hand, and the prediction of slot and intent labels, on the other hand.

6.2. Background noise

The VocADom@A4H test set was recorded in presence of background noise such as radio, television, etc. Some of these classes of noise have high frequency components, as for example those of vacuum cleaners. We randomly selected utterances that we annotated with background noise labels until we had about 10% (204 utterances) test set utterances with voice commands containing background noise available.

⁷<https://www.fon.hum.uva.nl/praat/>

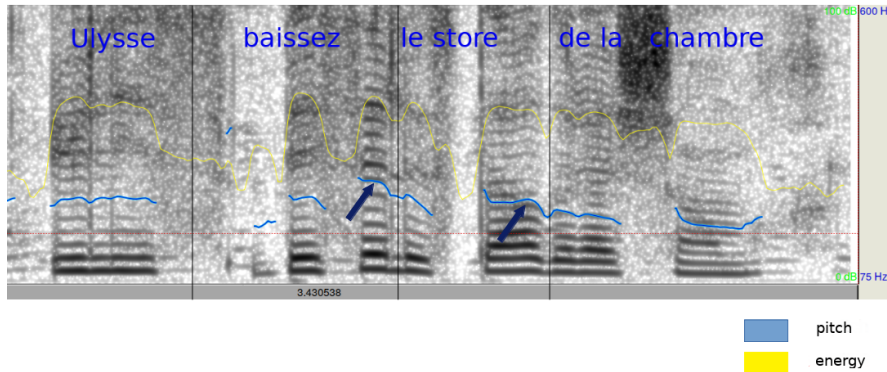


Figure 9: Pitch and energy for the utterance “ulyse lower the blind in the room”

645 Table 11 shows that E2E SLU outperforms (*All*, CER) pipeline SLU, especially for utterances with high-pitched vacuum cleaner background noise. Although the pipeline Kaldi ASR module outperforms ESPnet ASR in general, performance for both models is closer for utterances with vacuum cleaner background noise. This is the case for female speakers (*F*) as well as for male
650 speakers (*M*). These results can be related to the results of (Qian et al., 2016) according to which an E2E ASR system shows more robustness in processing noisy speech due to its layers of CNN (Convolutional Neural Network) in their architecture. However, here it is not the ASR task that benefits from it but the SLU task. It also seems that the Kaldi module exhibits much larger WER
655 variation between gender than ESPNet. However, this behavior is reversed for the CER (since the pipeline NLU is only processing words) suggesting that the E2E SLU is influenced by acoustic features late in the prediction process to exhibit such a gender bias.

6.3. Pitch and energy

660 Pitch and energy are known to impact ASR performances. In (Goldwater et al., 2010), the prosodic characteristics that are related to an increased ASR WER error have been studied and the authors concluded that *pitch* and intensity have an impact at extreme values. Pitch is a perceptive frequency-related

Table 11: ASR and SLU performances - voice commands with background noise

Background noise	Pipeline SLU		E2E SLU		# utt.
	WER (%)	CER (%)	WER (%)	CER (%)	
All:					
M&F	38.58	57.80	57.53	39.73	204
M	30.78	54.98	52.74	34.05	152
F	58.72	65.06	69.87	54.38	52
Vacuum cleaner:					
M&F	57.00	77.62	59.00	53.79	108
M	46.64	75.46	54.23	47.82	72
F	77.75	81.94	71.31	65.74	36
Radio&TV:					
M&F	20.31	35.77	56.31	25.94	75
M	18.08	36.64	53.30	20.90	58
F	27.96	32.84	66.57	43.13	17
Fan:					
M&F	14.74	32.69	65.36	26.92	21
M	13.18	24.99	62.84	38.88	17
F	18.27	49.99	71.03	0	4

account of a sound wave, which cannot be measured directly. However, two
665 tones can be considered to have the same pitch if they share the same F0 val-
ues. Regarding speech, increasing and decreasing pitch contours help define
prosody (Plack & Oxenham, 2005). According to (Stehwien & Vu, 2016) and
(Su & Tseng, 2018), most of the words related to concepts also carry a pitch
accent and can point to the most salient semantic information. In order to verify
670 this, they studied the correlation between pitch variations and concepts. These
studies inspired us to analyze the relationships between pitch, energy and SLU
performances.

The study was performed at utterance level on the test corpus VocADom@A4H. For each utterance, the ASR performance (WER) was computed for both pipeline and E2E models and the mean F0 and mean energy were computed using the *Praat* software.

Two F0 values were computed using two filters:

- A band-pass filter between 75 and 600Hz, typically containing male and female speaker F0 values and, on the other hand;
- No filter.

Hence for the F0 values without filters, high-pitched background noise values are also included. As the target of this study is the extraction of concepts and intents, the correlations were computed for the entire test set, as well as only for the 2612 test set utterances containing a voice command for comparison.

Table 12: Pearson and Spearman correlations energy and pitch with WER of pipeline ASR (Kaldi) and E2E ASR (ESPnet)

Correlation	Energy		Pitch (no filter)		Pitch(75-600Hz)	
	Kaldi	ESPnet	Kaldi	ESPnet	Kaldi	ESPnet
full dataset:						
Pearson (r)	-0.22**	-0.07**	-0.02	0.003	0.05**	0.007
Spearman (r_s)	-0.14**	-0.08**	-0.07**	-0.03*	0.01	0.002
voice commands only:						
Pearson (r)	0.04*	0.04*	0.23**	0.18**	0.09**	0.06*
Spearman (r_s)	0.05*	0.06**	0.10**	0.19**	0.06**	0.08**

* $p < 0.05$; ** $p < 0.01$

Table 12 shows that Word Error Rate is significantly correlated to energy but has a real impact only on Kaldi in the full data-set case. In all other cases, in particular for voice commands the correlation is negligible.

Regarding pitch, it is more correlated to WER for utterances with voice commands. This is especially the case for pitch values without filter. We further

690 analyzed this effect in three steps:

1. For the voice command utterances the F0 values between 75 and 600Hz every 0.01 seconds were computed.
2. Timestamps for the word boundaries of the reference and hypothesis transcripts, with symbolic concept labels were generated, applying a forced
695 alignment. Timestamps of F0 values on the one hand and timestamps of the word boundaries, on the other hand, were then aligned.
3. Finally, the number of reference concept labels with the highest F0 values in the utterance that were correctly predicted in the hypothesis predictions were computed for pipeline and E2E SLU.

700 The frequency list in Table 13 shows that 3 concept labels (**device**, **action**, **location-room**) are among the 10 most frequent concepts and words with the highest F0 value per utterance. 47.79% of all voice commands (1222/2557 voice commands) in the test set contain a concept consisting of words with the highest F0 values per utterance. It turns out that speakers, by talking to the home sys-
705 tem, make more effort in uttering commands and thus speak with an increased intonation, which results in higher F0 values for the words belonging to concepts of the command.

Table 13: Frequency of words and associated concepts with highest F0 value per utterance over 2557 voice commands of the test set

Frequency	Word	Frequency	Word
538	} (device)	62	vocadom
509	^ (action)	59	est-ce
163	cirrus	51	hé
160	dis	47	hestia
131	ulyse	43	chanticou
131	> (location-room)	39	allô
105	téraphim	37	que
84	ichéfix	35	messire
72	minouche	32	, (device-setting)

For the 1222 voice commands (reference utterances) with the highest concept pitch value, we calculated whether these concepts were well predicted by

710 the pipeline and E2E SLU systems. As shown in Table 14, a slightly higher percentage of concepts are retrieved for the E2E model as compared to the SLU pipeline model. Although the difference is small, this might indicate that the E2E SLU is slightly less impacted by pitch effects than the pipeline SLU.

Table 14: Reference voice command concepts with highest F0 values in hypothesis transcriptions

SLU model	concept ref. in hyp.(%)
Pipeline	74.22
E2E	75.00

6.4. Impact of MFCC and fbank features

715 Another way to check the impact of pitch on ASR and SLU performance is the removal of pitch variation from the test set utterances. To this end, average values of F0 per speaker were calculated. Using Praat, all test utterances for each speaker were resynthesized based on the resulting F0 average. As a next step, an E2E ASR model was also trained, using *MFCC* features, instead
 720 of *fbank* features, in order to compare the performance of ESPnet with the same acoustic features as used for Kaldi. Table 15 shows that, at the level of the pipeline SLU ASR module (*Pipeline ASR*), the performances of Kaldi (MFCC) on data *without* pitch variation are superior to performances on data *with* pitch variation. However *E2E ASR* and *E2E SLU* performances on data
 725 *with* pitch variation are superior to utterances *without* pitch variation, especially with fbank features.

Finally utterances from male ($M(1)$) and female ($F(1)$) speakers were evaluated separately. These utterances were compared with those of male ($M(2)$) and female ($F(2)$) speaker samples with a pitch *above* the average per speaker.
 730 These results show that performances for ($M(2)$) and ($F(2)$) with *deletion* of pitch variation are significantly worse than with pitch variation. This indicates that the E2E SLU is more robust to pitch variation.

Table 15: ASR and SLU performances (%), deletion of pitch variation

Model	Acoust. param.	No pitch var.		Pitch var.	
		WER	CER	WER	CER
Pipeline ASR:					
Kaldi	MFCC	21.48	-	22.92	-
E2E ASR:					
ESPNet ASR:	MFCC	49.90	-	47.60	-
	fbank	50.20	-	46.50	-
E2E SLU:	fbank	-	40.02	-	32.12
M(1).	fbank	-	41.94	-	32.90
F(1).	fbank	-	36.58	-	30.74
F0 > F0 avg.					
M(2).	fbank	-	53.32	-	42.40
F(2).	fbank	-	37.89	-	32.36

7. Grammatical Analysis of E2E SLU prediction

Although some pipeline approaches to SLU are able to handle OOV and syntactic variation, the E2E approach builds its own internal representation of utterances and has as final target a character string generation. This model should thus be more robust to OOV than classical ASR/NLU modules. To assess the capability of the E2E SLU model to handle better grammatical variations than a pipeline SLU system, we generated new input stimuli for which we controlled the *linguistic variation*, in particular at the *lexical* and *syntactic* levels. This is detailed in the two following sections.

7.1. Out of vocabulary words (OOV)

In order to measure the impact of an increased OOV rate, we gradually replaced the test set vocabulary of some specific concepts with words that did not appear in the training set. To measure an increasing difficulty this experiment was performed in 4 steps :

1. Step 1: action and device-setting,
2. Step 2: Step 1 and device ,
3. Step 3: Step 2 and location,
- 750 4. Step 4: Step 3 and keywords.

The following example shows a voice command (“vocadom turn on the kettle”) with symbolic intent and concepts *before* (1) and *after* (2) insertion of OOV words (Step 4):

(1) @ ah vocadom euh ^allume^ }la bouilloire} @
 755 (ah vocadom uh turn on the kettle)

(2) @ ah ursule euh ^enclenche^ }la bouillotte} @
 (ah ursule uh switch on the kettle)

Table 16 shows that substituted words in step 4 represent 26.15 % of the
 760 total number of word tokens (31k) and 3.48 % of the total number (1462) of word types (vocabulary).

Table 16: Vocadom@A4H - ratio OOV total words

Substitutions	#Word Type	#Words	(%) Word Type	(%) Total Words
Step 1	22	1785	1.50	5.72
Step 2	34	4276	2.32	13.70
Step 3	41	5516	2.80	17.68
Step 4	51	8160	3.48	26.15

Once the test set has been altered with OOV words, the speech utterances were generated with the same TTS tool as used for the artificial corpus generation in Section 4.2. The resulting utterances were then fed to the two SLU
 765 models. However, the E2E SLU was trained on data containing artificial speech, which was not the case for the pipeline SLU approach. Hence, for a fair comparison, the resulting E2E ASR model was used as ASR front end for the pipeline SLU.

Table 17 shows ASR and SLU performances according to the rate of OOV
 770 words. The complete *Compl. real.* line corresponds to the original Vocadom@A4H

Table 17: Impact of OOV on SLU performances (%)

Model	Pipeline				E2E		
	NLU		ASR+NLU		SLU		
	CER	F1	CER	F1	WER	CER	F1
Compl. real	33.78	85.51	36.24	84.21	46.50	32.12	74.57
Compl. synth.	-	-	37.07	83.34	39.30	25.00	53.70
OOV:							
Step 1	37.75	81.50	45.43	79.56	44.00	30.75	50.39
Step 2	53.77	72.39	62.03	72.48	53.20	46.75	50.26
Step 3	63.01	69.58	68.07	70.29	52.50	50.89	51.59
Step 4	90.45	63.66	86.44	65.03	55.90	58.80	51.43
Diff.	56.67	21.85	49.37	18.31	16.6	33.8	2.27

test set while the *Compl. synth.* lines correspond to the test set of which speech has been synthesized through TTS. Using the E2E ASR for the pipeline SLU did increase the CER (ASR+NLU column) but had a small impact on the intent prediction. Overall, for all models performances for concept (*CER*) and intent prediction (*F1*-score) deteriorate with increased OOV rates. The differences (*Diff.*) between *Compl. synth.* and *Step 4* are much smaller for the E2E model than for the pipeline SLU model for both concept and intent prediction. This would suggest that the E2E model is more robust to OOV words.

It must be noticed that the intent prediction of the E2E model, decreases dramatically with synthetic speech. This is due to an increased error rate for the *None* intent that consisted only of real speech in the training data whereas we used synthetic evaluation data for OOV impact evaluation. This is another evidence that acoustic features play a role up to the higher decision stages of the E2E model.

7.2. Syntactic variation

In this section, we measure the robustness of both SLU models for syntactic variability, predicting concepts and intents on test data with progressive

syntactic variability in two steps:

1. Step 1, 32 verbs belonging to the **action** concept were replaced by more complex syntactic constructions;
2. Step 2, substitutions in Step 1 have been augmented by disfluencies surrounding the words of 18 labelled concepts of **device**.

The following example ('Vocadom turn on the kettle') shows a voice command containing an intent and symbolic concepts from the test set *before* (1) and *after* (2) insertion of more complex syntactic constructions and disfluencies (Step 2):

- (1) @ vocadom euh ^allume^ }la bouilloire} @
(vocadom uh turn on the kettle)
- (2) @ vocadom euh pourrais-tu ^allumer^ la la }bouilloire} @
(vocadom uh could you turn on the the kettle)

We also generated text-to-speech based on the resulting modified test sets that we evaluated in the same way as for the OOV words test set up, as outlined in the previous section.

Table 18: Impact of syntactic variation on SLU performances (%)

Model	Pipeline				E2E		
	NLU		ASR+NLU		SLU		
	CER	F1	CER	F1	WER	CER	F1
Compl. real	33.78	85.51	36.24	84.21	46.50	32.12	74.57
Compl. synth.	-	-	37.07	83.34	39.30	25.00	53.70
Synt. var.:							
Step 1	38.41	81.06	50.40	77.45	44.40	16.29	52.59
Step 2	38.34	81.19	52.75	76.36	50.90	22.07	49.09
Diff.	4.56	4.32	15.68	6.98	11.60	2.93	4.61

Table 18 shows that differences (*Diff.*) between performances of concepts and intent prediction for test data with the complete original syntactic structure

Compl. synth. on the one hand, and test utterances with modified syntactic structure *Synt. var., Step 2* on the other hand, are smaller for the E2E model than for the sequential SLU model. This again indicates a greater robustness of the E2E model to cope with increased syntactic variation. Table 18 also shows that the performance of the E2E model for concept prediction improves with a more complex syntax. This may be due to an average sentence length of 15 words for the artificial corpus utterances, while the average sentence length for the (original) test set utterances is only 5. The increased syntactic variation, also increases the length of the evaluation utterances which consequently approaches the average length of the artificial corpus utterances.

Table 19: Impact of utterance length in the real test set case

Test set	Pipeline SLU		E2E SLU	
	Concept	Intent	Concept	Intent
	CER (%)	F1 (%)	CER (%)	F1 (%)
Compl. real 6747 utterances	36.24	84.21	32.12	74.57
Compl. real > 7 words 1461 utterances	38.62	80.01	33.18	68.23
Diff.	2.38	4.2	1.06	6.34

To test whether such robustness of the E2E SLU model is observable on the *real* test data (and not only on synthetic data), we extracted the more complex utterances from the test set based on their length. Since the average sentence length of the test set is seven words, 1461 utterances of the original test set (over 6747 total utterances) were labelled more complex. Although higher number of words does not necessarily mean more complex sentences, this is nevertheless correlated in the voice command case. The performances are shown on Table 19. The first row recalls the results on the whole test set while the second row shows performances computed on long utterances only. As with the synthetic dataset, the E2E SLU concept prediction error increases far less with the length

of sentences than the pipeline SLU (Table 19, Diff.). Furthermore, similarly to Table 9, we can see that the E2E SLU could not outperform the pipeline SLU intent prediction performances.

The reader can find a summary of the experiments measuring the acoustic
830 and grammatical impact on the E2E SLU model compared with the pipeline SLU model in Table B.20.

8. Discussion

To deal with the lack of data our strategy consisted in generating artificial utterances for voice commands. To deal with the bottleneck of distance between
835 real speech test data and artificial speech training data, we applied a *transfer learning* approach. An initial model was pre-trained on large out-of-domain data but composed of real speech data, and then this model was fine-tuned on the artificial but in-domain speech. It allowed us to take advantage of a large non-domain-specific data set, and a small domain specific data set. This approach
840 outperformed the pipeline SLU for concept prediction. Our data augmentation technique is close to (Li et al., 2018; Lugosch et al., 2020), who reported better concept prediction performances with a real speech model augmented with artificial speech than with an acoustic model only trained on real speech. Li et al. (2018) reported optimal performances for their E2E ASR model using
845 an acoustic model trained on 50% synthetic speech data and 50% real speech. On the other hand, *intent* prediction did not sufficiently benefit from transfer learning. A possible explanation is that ESPnet, being an ASR tool, functions at a level that is too local to be able to perform the global abstraction required for intent prediction.

850 Regarding the analysis of performances, for the Pipeline SLU model, performance differences between the NLU module and the complete sequential SLU system often remain high, despite the strategies used to reduce them. Since the current state of the art indicates that good SLU performance requires good ASR performance, we compared the pipeline and E2E SLU approaches on an

855 ASR task learned on equivalent data. Our experiments show that the WER of
the E2E ASR is significantly higher than the one of the pipeline ASR module.
However, the E2E SLU approach shows better SLU performance for concept
prediction, using the same tool (ESPnet) as E2E ASR. On top of that, our cor-
relation tests in Section 5.3 showed that perfect ASR is not necessary to obtain
860 good E2E SLU performance. It is, however, essential in the case of a pipeline
approach as we have demonstrated in (Desot et al., 2019b) for intent prediction
and in (Desot et al., 2019a) for concept prediction. **This answers our first
question and confirms the state-of-the-art: the E2E model reduces
the cascade of error effect.**

865 The E2E approach infers concepts and intents conveyed by an utterance
directly from the acoustic signal. Our experiments in Section 6 reveal that
prosodic information allows the model to point to the most important semantic
information. There are indications that the higher pitch values improve the
performance of the E2E SLU approach that turns out to be more robust to noisy
870 speech as compared to the pipeline model. This indicates that the convolutional
network of the E2E SLU model seems to benefit more from correlated and richer
filter-bank features than from MFCC features, used by the pipeline ASR model.
**This answer our second research question by showing that the E2E
SLU model uses prosodic information to infer concepts and is able to
875 learn a feature representation more robust to noise than the pipeline
model.**

As our target users are senior adults who tend to deviate easily from a fixed
grammar of voice commands, we tested the SLU approaches with increased
amount of OOV words and inserted more syntactic variation in the test corpus
880 utterances. In these two cases, the E2E SLU model proved to be more robust
than the pipeline approach while both models have been exposed to the same
data to learn the concepts to extract (same NLU training set). **This answers
our third research question. The E2E SLU model is indeed more
robust to syntactic variation than the pipeline model.**

885 9. Conclusion and future work

Our answer to the bottleneck of the cascade of error effect between the ASR and NLU modules of a pipeline SLU model, was an E2E SLU model that extracts intents and concepts directly from the signal. This approach based on deep neural networks allowed us to avoid the *cascade of errors* by performing a joint learning of these two tasks in one and the same model. By comparing our E2E SLU approach with a pipeline baseline approach, composed of a state-of-the-art ASR system and an NLU module, learned on data, specific to the home automation field, we were able to show that the E2E SLU approach gives best performance in terms of concept prediction. A possible solution to improve E2E SLU intent prediction is adding a decoder to the ESPnet architecture in order to train and predict concepts and intents jointly. A similar multi-task learning has already been applied for NLU in (Liu & Lane, 2016).

We can confirm one of the conclusions of the study of Stehwien & Vu (2016) and Su & Tseng (2018) that prosodic information can point to the most important semantic information. We have shown that the E2E SLU model exploits prosodic information which favours its performance in predicting intents and in particular concepts. On top of that the E2E SLU model shows more robustness as compared to the pipeline approach for processing target users syntactic variation and an increased OOV rate.

A transfer learning allowed us to decrease the *acoustic* distance between the artificial speech of the training data and the real speech of the test set. To further decrease this distance, speech synthesis based on a larger number of voices could be generated and added to the training data. Another possibility is training a neural speech synthesis model such as *Tacotron* (Wang et al., 2017; Li et al., 2018) for one or more speakers of the real SWEET-HOME corpus. The resulting model could then be used to generate new synthetic utterances that would be closer to the reference corpus.

One of the strengths of the E2E SLU approach is its processing of noisy speech data, which makes this approach very suitable in a realistic smart home

915 situation where voice commands must be extracted from utterances with various
background noises. In addition to background noise, residents are far away
from microphones. This leads to distorted acoustic signals by reverberation
depending on the acoustics of the room. To better assess and understand the
performance of the E2E SLU system in such a realistic situation, as a next
920 step, the micro-distant recording version of the test corpus should be tested.
These are the recordings made by the four antennas of 4 microphones integrated
into the ceiling of the Amiqua4Home smart home. This could be solved by
augmenting our training data with *room impulse response* (RIR) data. However,
the acquisition of real RIR data is not trivial. Ko et al. (2017) show that
925 acoustic models trained on simulated RIR data are competitive with real RIR
data. Therefore, a technique for increasing our training data with simulated
RIR data could be explored.

Acknowledgements

This work is part of the VOCADOM project founded by the French National
930 Research Agency (Agence Nationale de laRecherche) / ANR-16-CE33-0006. It
was also partially supported by MIAI@Grenoble-Alpes (ANR-19-P3IA-0003).

References

- Abdel-Hamid, O., Mohamed, A.-R., Jiang, H., & Penn, G. (2012). Applying
convolutional neural networks concepts to hybrid NN-HMM model for speech
935 recognition. In *International Conference on Acoustics, Speech, and Signal
Processing (ICASSP)* (pp. 4277–4280). IEEE.
- Amodei, D., Ananthanarayanan, S., Anubhai, R., Bai, J., Battenberg, E., Case,
C., Casper, J., Catanzaro, B., Cheng, Q., Chen, G. et al. (2016). Deep
speech2: End-to-end speech recognition in English and Mandarin. In *Inter-*
940 *national conference on machine learning (ICML)* (pp. 173–182).

- Anastasakos, T., McDonough, J., Schwartz, R., & Makhoul, J. (1996). A compact model for speaker-adaptive training. In *International Conference on Spoken Language Processing (ICSLP)* (pp. 1137–1140). IEEE volume 2.
- Bahdanau, D., Cho, K., & Bengio, Y. (2014). Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, .
- Bapna, A., Tür, G., Hakkani-Tür, D., & Heck, L. (2017). Sequential Dialogue Context Modeling for Spoken Language Understanding. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue* (pp. 103–114). Saarbrücken, Germany: Association for Computational Linguistics.
- 945 Bird, S., Klein, E., & Loper, E. (2009). *Natural language processing with Python: analyzing text with the natural language toolkit*. ” O’Reilly Media, Inc.”.
- Caubrière, A., Rosset, S., Estève, Y., Laurent, A., & Morin, E. (2020). Where are we in named entity recognition from speech? In *Proceedings of The 12th Language Resources and Evaluation Conference* (pp. 4514–4520).
- 955 Caubrière, A., Tomashenko, N., Laurent, A., Morin, E., Camelin, N., & Estève, Y. (2019). Curriculum-based transfer learning for an effective end-to-end spoken language understanding and domain portability. In *interspeech* (pp. 1198–1202).
- Cho, J., Baskar, M. K., Li, R., Wiesner, M., Mallidi, S. H., Yalta, N., Karafiat, 960 M., Watanabe, S., & Hori, T. (2018). Multilingual sequence-to-sequence speech recognition: architecture, transfer learning, and language modeling. In *IEEE Spoken Language Technology Workshop (SLT)* (pp. 521–527). IEEE.
- Crystal, D. (2011). *A Dictionary of Linguistics and Phonetics*. The Language Library. Wiley.
- 965 Denisov, P., & Vu, N. T. (2020). Pretrained semantic speech embeddings for end-to-end spoken language understanding via cross-modal teacher-student learning. *arXiv preprint arXiv:2007.01836*, .

- Desot, T., Portet, F., & Vacher, M. (2019a). SLU for voice command in smart home: comparison of pipeline and end-to-end approaches. In *Automatic Speech Recognition and Understanding Workshop (ASRU)* (pp. 822–829). IEEE.
- Desot, T., Portet, F., & Vacher, M. (2019b). Towards end-to-end spoken intent recognition in smart home. In *Conference on Speech Technology and Human-Computer Dialogue (SpeD)* (pp. 1–8).
- Desot, T., Portet, F., & Vacher, M. (2020). Corpus generation for voice command in smart home and the effect of speech synthesis on end-to-end SLU. In *International Conference on Language Resources and Evaluation (LREC)* (pp. 6395–6404).
- Desot, T., Raimondo, S., Mishakova, A., Portet, F., & Vacher, M. (2018). Towards a French smart-home voice command corpus: Design and NLU experiments. In *International Conference on Text, Speech, and Dialogue (TSD)* (pp. 509–517). Springer.
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)* (pp. 4171–4186). Association for Computational Linguistics.
- Digalakis, V. V., & Neumeyer, L. G. (1996). Speaker adaptation using combined transformation and bayesian methods. *IEEE transactions on speech and audio processing*, 4, 294–300.
- Galliano, S., Geoffrois, E., Mostefa, D., Choukri, K., Bonastre, J.-F., & Gravier, G. (2005). The ESTER phase ii evaluation campaign for the rich transcription of French broadcast news. In *European Conference on Speech Communication and Technology (EUROSPEECH)*.

- Galliano, S., Gravier, G., & Chaubard, L. (2009). The ESTER 2 evaluation campaign for the rich transcription of French radio broadcasts. In *Tenth Annual Conference of the International Speech Communication Association*.
- Gatt, A., & Krahmer, E. (2018). Survey of the state of the art in natural language generation: Core tasks, applications and evaluation. *Journal of Artificial Intelligence Research*, 61.
- Ghannay, S., Caubrière, A., Estève, Y., Laurent, A., & Morin, E. (2018). End-to-end named entity extraction from speech. *arXiv preprint arXiv:1805.12045*, .
- Giraudel, A., Carré, M., Mapelli, V., Kahn, J., Galibert, O., & Quintard, L. (2012). The REPERE Corpus: a multimodal corpus for person recognition. In *International Conference on Language Resources and Evaluation (LREC)* (pp. 1102–1107).
- Goldwater, S., Jurafsky, D., & Manning, C. D. (2010). Which words are hard to recognize? Prosodic, lexical, and disfluency factors that increase speech recognition error rates. *Speech Communication*, 52, 181–200.
- Graves, A., Fernández, S., Gomez, F., & Schmidhuber, J. (2006). Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *International conference on Machine learning (ICML)* (pp. 369–376). ACM.
- Gravier, G., Adda, G., Paulson, N., Carré, M., Giraudel, A., & Galibert, O. (2012). The ETAPE corpus for the evaluation of speech-based TV content processing in the French language. In *International Conference on Language Resources and Evaluation (LREC)*.
- Hahn, S., Lehnen, P., Raymond, C., & Ney, H. (2008). A comparison of various methods for concept tagging for spoken language understanding. In *International Conference on Language Resources and Evaluation (LREC)*.

- Hakkani-Tür, D., Béchet, F., Riccardi, G., & Tur, G. (2006). Beyond ASR 1-best: Using word confusion networks in spoken language understanding. *Computer Speech & Language*, *20*, 495–514.
- 1025
- Hannun, A., Case, C., Casper, J., Catanzaro, B., Diamos, G., Elsen, E., Prenger, R., Satheesh, S., Sengupta, S., Coates, A. et al. (2014). Deep speech: Scaling up end-to-end speech recognition. *arXiv preprint arXiv:1412.5567*, .
- He, Y., & Young, S. (2003). A data-driven spoken language understanding system. In *Workshop on Automatic Speech Recognition and Understanding (ASRU)* (pp. 583–588). IEEE.
- 1030
- Hemphill, C. T., Godfrey, J. J., & Doddington, G. R. (1990). The ATIS spoken language systems pilot corpus. In *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27,1990*.
- Hou, Y., Liu, Y., Che, W., & Liu, T. (2018). Sequence-to-sequence data augmentation for dialogue language understanding. In *Proceedings of the 27th International Conference on Computational Linguistics (COLING)* (pp. 1234–1245).
- 1035
- Huang, L., Sil, A., Ji, H., & Florian, R. (2017). Improving slot filling performance with attentive neural networks on dependency structures. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)* (pp. 2588–2597).
- 1040
- Jeong, M., & Lee, G. G. (2008). Triangular-chain conditional random fields. *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, *16*, 1287–1302. doi:10.1109/TASL.2008.925143.
- 1045
- Ko, T., Peddinti, V., Povey, D., Seltzer, M. L., & Khudanpur, S. (2017). A study on data augmentation of reverberant speech for robust speech recognition. In *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)* (pp. 5220–5224). IEEE.

- 1050 Krueger, K. A., & Dayan, P. (2009). Flexible shaping: How learning in small steps helps. *Cognition*, *110*, 380–394.
- Leggetter, C., & Woodland, P. (1995). Maximum likelihood linear regression for speaker adaptation of continuous density hidden Markov models. *Computer Speech & Language*, *9*, 171 – 185.
- 1055 Li, J., Gadde, R., Ginsburg, B., & Lavrukhin, V. (2018). Training neural speech recognition systems with synthetic speech augmentation. *arXiv preprint arXiv:1811.00707*, .
- Liu, B., & Lane, I. (2016). Attention-based recurrent neural network models for joint intent detection and slot filling. In *Interspeech* (pp. 685–
1060 689). URL: http://www.isca-speech.org/archive/Interspeech_2016/abstracts/1352.html. doi:10.21437/Interspeech.2016-1352.
- Liu, C., Zhu, S., Zhao, Z., Cao, R., Chen, L., & Yu, K. (2020). Jointly encoding word confusion network and dialogue context with bert for spoken language understanding. In *Interspeech* (pp. 871–875).
- 1065 Lugosch, L., Meyer, B., Nowrouzezahrai, D., & Ravanelli, M. (2020). Using speech synthesis to train end-to-end spoken language understanding models. In *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. IEEE.
- Lugosch, L., Ravanelli, M., Ignoto, P., Tomar, V. S., & Bengio, Y. (2019).
1070 Speech model pre-training for end-to-end spoken language understanding. In *Interspeech* (pp. 814–818).
- Mangu, L., Brill, E., & Stolcke, A. (2000). Finding consensus in speech recognition: word error minimization and other applications of confusion networks. *Computer Speech & Language*, *14*, 373 – 400.
- 1075 Mesnil, G., Dauphin, Y., Yao, K., Bengio, Y., Deng, L., Hakkani-Tur, D., He, X., Heck, L., Tur, G., Yu, D., & others (2015). Using recurrent neural networks

- for slot filling in spoken language understanding. *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, 23, 530–539. URL: <http://dl.acm.org/citation.cfm?id=2876380>.
- 1080 Mishakova, A., Portet, F., Desot, T., & Vacher, M. (2019). Learning natural language understanding systems from unaligned labels for voice command in smart homes. In *IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)* (pp. 832–837).
- 1085 Möller, S., Gödde, F., & Wolters, M. (2008). Corpus analysis of spoken smart-home interactions with older users. In *Proceedings of the 6th International Conference on Language Resources and Evaluation*.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., & Lerer, A. (2017). Automatic differentiation in Pytorch. In *Advances in Neural Information Processing Systems (NIPS) workshop*.
- 1090
- Plack, C. J., & Oxenham, A. J. (2005). Overview: The present and future of pitch. In *Pitch* (pp. 1–6). Springer.
- Portet, F., Caffiau, S., Ringeval, F., Vacher, M., Bonnefond, N., Rossato, S., Lecouteux, B., & Desot, T. (2019). Context-Aware Voice-based Interaction in Smart Home-VocADom@ A4H Corpus Collection and Empirical Assessment of its Usefulness. In *2019 IEEE Intl Conf on Dependable, Autonomic and Secure Computing, Intl Conf on Pervasive Intelligence and Computing, Intl Conf on Cloud and Big Data Computing, Intl Conf on Cyber Science and Technology Congress (DASC/PiCom/CBDCCom/CyberSciTech)* (pp. 811–818). IEEE.
- 1095
- 1100 Povey, D., Ghoshal, A., Boulianne, G., Burget, L., Glembek, O., Goel, N., Hannemann, M., Motlicek, P., Qian, Y., Schwarz, P. et al. (2011). The Kaldi speech recognition toolkit. In *Workshop on Automatic Speech Recognition and Understanding (ASRU)*.

- 1105 Povey, D., Zhang, X., & Khudanpur, S. (2015). Parallel training of DNNs with natural gradient and parameter averaging. In *International Conference on Learning Representations (ICLR)*.
- Qian, Y., Bi, M., Tan, T., & Yu, K. (2016). Very deep convolutional neural networks for noise robust speech recognition. *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, *24*, 2263–2276.
1110
- Qian, Y., Ubale, R., Ramanaryanan, V., Lange, P., Suendermann-Oeft, D., Evanini, K., & Tsuprun, E. (2017). Exploring ASR-free end-to-end modeling to improve spoken language understanding in a cloud-based dialog system. In *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)* (pp. 569–576). IEEE.
1115
- Rao, M., Raju, A., Dheram, P., Bui, B., & Rastrow, A. (2020). Speech to semantics: Improve ASR and NLU jointly via all-neural interfaces. *arXiv preprint arXiv:2008.06173*, .
- Sears, J. A. (1988). The DARPA spoken language systems program: Past, present, and future. *The Journal of the Acoustical Society of America*, *84*, S188–S188.
1120
- Serdyuk, D., Wang, Y., Fuegen, C., Kumar, A., Liu, B., & Bengio, Y. (2018). Towards end-to-end spoken language understanding. In *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)* (pp. 5754–5758).
1125 IEEE.
- Serpollet, N., Bergounioux, G., Chesneau, A., & Walter, R. (2007). A large reference corpus for spoken French: ESLO 1 and 2 and its variations. In *Proceedings from Corpus Linguistics Conference Series, University of Birmingham*.
- 1130 Simonnet, E., Ghannay, S., Camelin, N., Estève, Y., & De Mori, R. (2017). ASR error management for improving spoken language understanding. In *Interspeech* (pp. 3329–3333).

- Stehwien, S., & Vu, N. T. (2016). Exploring the correlation of pitch accents and semantic slots for spoken language understanding. In *Interspeech* (pp. 730–734).
1135
- Su, C.-Y., & Tseng, C.-Y. (2018). Perceivable information structure in discourse prosody-detecting prominent prosodic words in spoken discourse using F0 contour. In *2018 11th International Symposium on Chinese Spoken Language Processing (ISCSLP)* (pp. 424–428). IEEE.
- 1140 Sudoh, K., Tsukada, H., & Isozaki, H. (2006). Incorporating speech recognition confidence into discriminative named entity recognition of speech data. In *International Conference on Computational Linguistics* (pp. 617–624). Association for Computational Linguistics.
- Takahashi, S.-y., Morimoto, T., Maeda, S., & Tsuruta, N. (2003). Dialogue
1145 experiment for elderly people in home health care system. In *International Conference on Text, Speech, and Dialogue (TSD)* (pp. 418–423).
- Tan, T.-P., & Besacier, L. (2006). A French Non-Native Corpus for Automatic Speech Recognition. In *International Conference on Language Resources and Evaluation (LREC)* (pp. 1610–1613). volume 6.
- 1150 Tokui, S., Oono, K., Hido, S., & Clayton, J. (2015). Chainer: a next-generation open source framework for deep learning. In *Proceedings of workshop on machine learning systems (LearningSys) in the twenty-ninth annual conference on Neural Information Processing Systems (NIPS)* (pp. 1–6). volume 5.
- Tur, G., & De Mori, R. (2011). *Spoken Language Understanding Systems for
1155 Extracting Semantic Information from Speech*. Wiley.
- Ueno, S., Inaguma, H., Mimura, M., & Kawahara, T. (2018). Acoustic-to-word attention-based model complemented with character-level CTC-based model. In *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)* (pp. 5804–5808). IEEE.

- 1160 Vacher, M., Bouakaz, S., Chaumon, M.-E. B., Aman, F., Khan, R. A., Bekkadj, S., Portet, F., Guillou, E., Rossato, S., & Lecouteux, B. (2016). The CIRDO corpus: comprehensive audio/video database of domestic falls of elderly people. In *International Conference on Language Resources and Evaluation (LREC)* (pp. 1389–1396).
- 1165 Vacher, M., Caffiau, S., Portet, F., Meillon, B., Roux, C., Elias, E., Lecouteux, B., & Chahuara, P. (2015). Evaluation of a context-aware voice interface for Ambient Assisted Living: qualitative user study vs. quantitative system evaluation. *ACM Transactions on Accessible Computing (TACCESS)*, 7, 5:1–5:36. doi:10.1145/2738047.
- 1170 Vacher, M., Fleury, A., Serignat, J.-F., Noury, N., & Glasson, H. (2008). Preliminary evaluation of speech/sound recognition for telemedicine application in a real environment. In *Interspeech* (pp. 496–499).
- Vacher, M., Lecouteux, B., Chahuara, P., Portet, F., Meillon, B., & Bonnefond, N. (2014). The Sweet-Home speech and multimodal corpus for home automation interaction. In *International Conference on Language Resources and Evaluation (LREC)* (pp. 4499–4506).
- 1175
- Wang, Y., Skerry-Ryan, R., Stanton, D., Wu, Y., Weiss, R. J., Jaitly, N., Yang, Z., Xiao, Y., Chen, Z., Bengio, S. et al. (2017). Tacotron: Towards end-to-end speech synthesis. In *Interspeech* (pp. 4006–4010).
- 1180 Watanabe, S., Hori, T., Karita, S., Hayashi, T., Nishitoba, J., Unno, Y., Soplín, N., Heymann, J., Wiesner, M., Chen, N., Renduchintala, A., & Ochiai, T. (2018). ESPNet: End-to-end speech processing toolkit. In *Interspeech* (pp. 2207–2211).
- 1185 Watanabe, S., Hori, T., Kim, S., Hershey, J. R., & Hayashi, T. (2017). Hybrid CTC/attention architecture for end-to-end speech recognition. *IEEE Journal of Selected Topics in Signal Processing*, 11, 1240–1253.

1190 Zhai, L., Fung, P., Schwartz, R., Carpuat, M., & Wu, D. (2004). Using n-best lists for named entity recognition from chinese speech. In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (HLT-NAACL)* (pp. 37–40). Association for Computational Linguistics.

Zhang, Y., Pezeshki, M., Brakel, P., Zhang, S., Bengio, C. L. Y., & Courville, A. (2016). Towards End-to-End Speech Recognition with Deep Convolutional Neural Networks. In *Interspeech 2016* (pp. 410–414).

1195 **Appendix A. Some examples of variants of a window opening command**

	French	English translation
1	Ouvre la fenêtre	Open the window
2	Ouvre la fenêtre s'il vous plaît	Open the window please
3	Est-ce que tu peux ouvrir la fenêtre?	Can you open the window?
4	Est-ce que tu peux ouvrir la fenêtre s'il vous plaît?	Can you open the window please?
5	Je veux que tu ouvres la fenêtre	I want you to open the window

Appendix B. Overview of measures of the acoustic and grammatical impact on the SLU models

Table B.20: Overview of Pipeline and E2E SLU evaluation results on VocADom@A4H test data

Analysis level:	Acoustic										Symbolic			
	Source					Prosody					Lexical			
	Noise		Gender			Avg. F0		MFCC			OOV			
V	RT	F	A	M	F	M&F	0	1	2	3	4	1	2	
Pipeline SLU :														
WER	57.00	20.31	14.74	38.58	22.25	24.12	22.92	21.48	-	-	-	-	-	-
Concept Error Rate	77.62	35.77	32.69	57.80	37.56	33.77	36.24	-	-	37.07	45.43	62.03	68.07	86.44
Intent F-measure	-	-	-	-	83.64	86.48	84.21	-	-	83.34	79.56	72.48	70.29	65.03
E2E SLU :														
WER	59.00	56.31	65.36	57.53	46.90	44.00	46.50	49.90	50.20	39.30	44.00	53.20	52.50	55.90
Concept Error Rate	53.79	25.94	26.92	39.73	32.90	30.71	32.12	-	40.02	25.00	30.75	46.75	50.89	58.80
Intent F-measure	-	-	-	-	74.30	75.04	74.57	-	-	53.70	50.39	50.26	51.59	51.43

Table B.21: Explication of terminology used in table B.20 summarizing SLU performances

Analysis	Explanation
Noise	V Vacuum cleaner RT Radio-Television F Fan A All background noises
Gender	M Masculine F Feminine M&F Masculine and Feminine
Avg.F0	mfcc, fbank Average F0 value per speaker, mfcc, fbank
Lexical OOV	0 Artificial Speech synthesis, complete vocabulary 1, 2, 3, 4 Out of vocabulary, steps 1, 2, 3, 4
Syntactic variation	1, 2 Steps 1, 2

- 1st paper studying acoustic and linguistic features used by E2E DNN models for SLU
- joint learning of ASR and NLU in E2E SLU does fight the cascade of errors effect.
- E2E SLU models exploit prosodic information to predict intents and concepts.
- E2E models are more robust than pipeline models to noise and grammatical variation
- transfer learning and artificial training data are efficient for low resource tasks

Declaration of interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: