



HAL
open science

Toward Low-Cost End-to-End Spoken Language Understanding

Marco Dinarelli, Marco Naguib, François Portet

► **To cite this version:**

Marco Dinarelli, Marco Naguib, François Portet. Toward Low-Cost End-to-End Spoken Language Understanding. Interspeech 2022, Sep 2022, Incheon, South Korea. pp.2728-2732, 10.21437/Interspeech.2022-10702 . hal-03872546

HAL Id: hal-03872546

<https://hal.science/hal-03872546v1>

Submitted on 6 Jan 2025

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.

Toward Low-Cost End-to-End Spoken Language Understanding

Marco Dinarelli, Marco Naguib, François Portet

Univ. Grenoble Alpes, CNRS, Grenoble INP, LIG, 38000 Grenoble, France

(marco.dinarelli|francois.portet)@univ-grenoble-alpes.fr, marco.naguib@hotmail.com

Abstract

Recent advances in spoken language understanding benefited from Self-Supervised models trained on large speech corpora. For French, the LeBenchmark project has made such models available and has led to impressive progress on several tasks including spoken language understanding. These advances have a non-negligible cost in terms of computation time and energy consumption. In this paper, we compare several learning strategies trying to reduce such cost while keeping competitive performance. At the same time we propose an extensive analysis where we measure the cost of our models in terms of training time and electric energy consumption, hopefully promoting a comprehensive evaluation procedure. The experiments are performed on the FSC and MEDIA corpora, and show that it is possible to reduce the learning cost while maintaining state-of-the-art performance and using SSL models.

Index Terms: speech recognition, human-computer interaction, computational paralinguistics

1. Introduction

Spoken Language Understanding (SLU) aims at extracting semantic representations from speech signals containing sentences in natural language [1]. Classical approaches to SLU used a cascade model made of an Automatic Speech Recognizer (ASR) feeding a Natural Language Understanding (NLU) module [2, 3, 4, 5]. Neural networks led to large advances for *end-to-end* SLU systems [6, 7, 8, 9, 10, 11], which are preferred to cascade systems, in particular for their ability to reduce error propagation effects and to exploit acoustic components to deduct semantic information [12].

Though some approaches proposed full end-to-end learning of SLU models, e.g. [13, 7, 14], most proposals use incremental learning on more and more specific tasks (e.g. ASR) before SLU. Based on the assumption that an SLU model must learn the lexical representation supporting meaning, [8] and [10] propose a model trained progressively to transcribe the signal and then to extract the semantics. In this trend there are also [15] and [16], which learn first a domain classifier, and then optimize the model for predicting user intent and semantic slots. Recently end-to-end SLU approaches integrated in the system Self-Supervised Learning (SSL) models trained on huge amount of speech data, e.g. *wav2vec* and *HuBERT* [17, 18, 19]. [20] for example, uses a pre-trained *wav2vec* model as speech encoder, while the *SUPERB* benchmark [21] proposes slot-filling and intent classification tasks to evaluate the SSL pre-trained models. SSL models are now available also for the French community [22, 23], and led to impressive improvements on downstream tasks, including SLU.

While such advances are a great scientific achievement, they come at non-negligible costs in terms of resource consumption [24, 25, 26]. For instance the largest model in [23] is trained during two weeks on 64 GPUs, and this is even not that expensive

compared to larger models or models trained on far more data [27]. We can argue this is still a limited cost as such models are trained once and used for many applications. However, SSL models are often a component of the whole system, and they need to be fine-tuned on the downstream task to achieve good results. This practice multiplies costly training phases, which may lead to huge resource consumption.

In this paper we aim at analysing and reducing the cost for obtaining competitive performance on SLU tasks while exploiting SSL pre-trained models such as those deployed by [22, 23]. We propose an analysis for finding a better compromise between performance and energy consumption. We show the computational cost of our models in terms of different measures, including training time, energy consumption and carbon footprint. Our analysis is based on different, yet simple training strategies with respect to those used in [22, 23]. In addition, for the French task MEDIA [28], we couple our training strategies with: *i*) a transfer learning strategy using models for a different task [29]; *ii*) an offline fine-tuning phase of a SSL model directly on the SLU target task, instead of the ASR task as proposed in [23]. Such fine-tuning is also performed for the FSC task [8] using the *XLSR53* SSL model [27]. While our fine-tuning is relatively expensive with respect to the downstream SLU model learning, it is still cheaper than recent approaches, while allowing for comparable performance.

2. SLU and SSL Models

In this paper we exploit pre-trained SSL models and we seek to get the best outcome as possible on SLU tasks using the least amount of resources. We use the SSL *w2v2-fr-7k* model from *LeBenchmark* for French [23], which led to the best performance on SLU. We use the multi-language *XLSR-53* model for English [27].¹

SLU models used in this paper are the same as in [23].² They are *sequence-to-sequence* models based on LSTMs and attention mechanisms [32, 33]. The encoder has a similar pyramidal structure as the one proposed in [34], the decoder uses two attention mechanisms, one for attending the encoder’s hidden states, and one for attending the decoder’s previous predictions, like the self-attention module of Transformers [35]. All models are trained minimizing the CTC loss [36]. In this work we use SSL models as feature extractors. Features are given as input to SLU models as an alternative to traditional features (e.g. MFCC). This usage enables a cheaper training phase than including the SSL model as a component of the SLU system.³

¹We were not able to use the English monolingual SSL model in our Fairseq architecture.

²We downloaded models and systems from <https://huggingface.co/LeBenchmark> and <https://github.com/LeBenchmark/NeurIPS2021> respectively.

³Comparing at the same effort, that is: offline feature extraction without fine-tuning vs. SSL model as frozen component of the system; or offline fine-tuning and feature extraction vs. SSL model as fine-tuned component of the system.

Corpus: PortMEDIA, Metric: Concept Error Rate (CER ↓)							
Train stg	# params (M)	Input	kWh (gCO ₂)	kWh/p	Train T	DEV	TEST
Base features							
3 steps	9.83	spectro	4,473 (228)	0,099	36h14'	35.91	40.57
2 steps	9.83	spectro	2,989 (152)	∞	24h14'	65.80	87.32
1 step	9.83	spectro	1,708 (87)	\mathcal{M}_2	15h52'	59.22	68.50
3 steps	11.76	w2v2-fr	3,983 (203)	2,235	36h22'	22.17	22.51
2 steps	11.76	w2v2-fr	2,707 (138)	1,939	24h27'	21.86	23.02
1 step	11.76	w2v2-fr	1,815 (93)	\mathcal{M}_2	18h08'	25.53	23.48
Fine-tuned features (+100h x4 GPU ⇒ 92.720 kWh / 4729 gCO ₂)							
1 step +1	11.76 (+318)	w2v2-fr slu	1,214 (62)	-	11h34'	21.50	22.13
Literature							
[30]	-	MFCC	-	-	-	-	42.3

Table 1: Results on the French corpus PortMEDIA. See the text for details.

Corpus: FSC, Metric: Accuracy ↑							
Train stg	# params (M)	Input	kWh (CO ₂)	kWh/p	Train T	DEV	TEST
Base features							
3 steps	7.93	spectro	1.474 (75)	0.259	12h32'	82.14	93.78
2 steps	7.93	spectro	0.707 (36)	0.208	5h57'	78.19	91.01
1 step	7.93	spectro	0.503 (26)	\mathcal{M}_2	5h02'	79.79	90.03
3 steps	9.70	XLSR-53	1.068 (54)	0.519	9h57'	97.63	99.66
2 steps	9.70	XLSR-53	0.950 (48)	0.564	8h34'	96.57	99.39
1 step	9.70	XLSR-53	0.668 (34)	\mathcal{M}_2	8h08'	96.66	98.89
Fine-tuned features (+100h x4 GPU ⇒ 92.720 kWh / 4729 gCO ₂)							
1 step +1	9.70 (+315)	XLSR-53 slu	0.365 (19)	-	3h23'	99.20	99.71
State of the art							
[8]	-	raw audio	-	-	-	-	98.70
[31]	287(*)	speech+text	-	-	-	-	99.71

Table 2: Results on the English FSC corpus. See the text for details. (*): our own estimation.

Models described in [23] are learned with three training steps. Each training step uses the model learned at the previous step for initializing the current model’s parameters. This strategy is named *3 steps* in this paper, and it consists in: **1)** training the encoder for ASR; **2)** training the encoder for SLU; **3)** training the whole model, encoder and decoder, for SLU. While this strategy is the most effective, it requires a high training cost. Additionally the best models proposed in [23] require an additional training step for supervised fine-tuning of the SSL model on the downstream task.

Results of SLU models using input features generated with SSL models are very high, we argue that a comparable performance can be reached with a lower training cost, reducing resource consumption. With our analyses we would like to reach a better compromise between training cost and model’s final performance. In order to validate our hypothesis we propose two alternative and simple training strategies (indicated as **Train stg** in our tables): in the “*2 steps*” strategy we perform only steps 2 and 3 of the *3 steps* strategy; in “*1 step*” we train directly the final SLU model full end-to-end.

3. Evaluation

3.1. Data

We will use mainly the MEDIA corpus for French [28], which we extensively used in the past [37, 38, 39, 40, 41, 42, 43], used also in [23] and allowing thus for direct comparison, and the FSC corpus for English [8].

The MEDIA corpus focuses on the domain of hotel information and reservation in France. It is made of 1,250 human-machine dialogues transcribed and annotated with 76 semantic concepts. The corpus is split into 12,908 utterances (41.5h) for training, 1,259 for development (3.5h), and 3,005 for test (11.3h). For French, in this paper we consider the MEDIA task as the target

task, that is the task on which we would like to reach the best possible performance at the lowest possible cost, starting possibly from existing resources such as pre-trained SSL models, but also pre-trained SLU models for other (possibly similar) tasks. In order to reach these experimental conditions and inspired by transfer learning techniques of [30], we use also the PORT-MEDIA (PM) corpus [29], focusing on ticket reservation for the *2010 Avignon Festival*. This corpus has been collected and annotated following the same paradigm as for MEDIA. PM is also divided in training, development and test splits, composed respectively of 5,900, 1,400 and 2,800 utterances. This corpus has been annotated with 36 semantic concepts close to the MEDIA concept set: PORTMEDIA and MEDIA share indeed 26 concepts.

In order to have a wider view on SLU results and their cost, we also used the *Fluent Speech Commands* (FSC) corpus [8], composed of 30,043 utterances from 97 speakers. Each utterance contains a vocal command for a domestic environment (e.g. *increase heat in the kitchen*), and is annotated with 3 attributes: action, object and localisation. In the original paper, authors define user intent as the combination of these 3 attributes, and SLU results are computed as accuracy on this combination. The corpus is split in 23,132 utterances for training (14.7h), 3,118 for development (1.9h), and 3,793 for test (2.4h). Since results on FSC are particularly high since its release [8], we did not perform additional experiments with transfer learning on this task.

For more details on data sets and experimental settings we refer the reader to previous work [23, 8].

3.2. Evaluating the computational cost

We evaluate the computational cost of our models measuring: the training time, the electric energy consumption in *kWh* and its conversion into grams of CO₂ (indicated as gCO₂ in the tables).

These 2 values are obtained using the *codecarbon* tool⁴. Since this tool overestimates the conversion between *kWh* and *gCO₂*, like in [26] we use the official coefficient of 51 grams/*kWh*.⁵ We also show the cost in *kWh* for gaining one point of the evaluation metric, either Concept Error Rate (CER) or Accuracy (Acc.). This metric is indicated as **kWh/p** in the tables. For computing this value we assume a cheaper model \mathcal{M}_2 , in terms of resource consumption, obtains worse results than a compared model \mathcal{M}_1 . Defining then $\text{kWh}(\mathcal{M}_i)$ and $\text{CER}(\mathcal{M}_i)$ respectively the energy consumption and the CER of the model \mathcal{M}_i , the value kWh/p is defined as $\frac{\text{kWh}(\mathcal{M}_1) - \text{kWh}(\mathcal{M}_2)}{\text{CER}(\mathcal{M}_2) - \text{CER}(\mathcal{M}_1)}$, with the constraint $\text{kWh}(\mathcal{M}_1) \geq \text{kWh}(\mathcal{M}_2)$. Since \mathcal{M}_2 is cheaper, the numerator is positive. Assuming \mathcal{M}_2 makes a greater amount of errors than \mathcal{M}_1 , the denominator is also positive. Defining analogously $\text{Acc}(\mathcal{M}_i)$ the accuracy of model \mathcal{M}_i , we can define kWh/p for tasks evaluated through accuracy, like FSC in this work: $\frac{\text{kWh}(\mathcal{M}_1) - \text{kWh}(\mathcal{M}_2)}{\text{Acc}(\mathcal{M}_1) - \text{Acc}(\mathcal{M}_2)}$, again we assume a cheaper model is less effective, and we swapped \mathcal{M}_1 and \mathcal{M}_2 at the denominator since the higher the accuracy the better. In the following tables, the value kWh/p is given with respect to the cheapest model using the same input features (**Input** in the tables). Such a model is indicated with \mathcal{M}_2 in the same column of the table. When a model \mathcal{M}_1 is worse both in terms of CER (or accuracy) and *kWh* than the reference model \mathcal{M}_2 , we use by convention the value ∞ for kWh/p , meaning that a more expensive training would not lead to any performance improvement. The interpretation of kWh/p should be understood keeping in mind the assumption that a cheaper model is less effective. Then, with given input features, kWh/p measures the additional cost in *kWh* to increase the performance by 1 point. Such cost could be due to using more data and/or a bigger model, leading to a longer training time. The kWh/p metric can be interpreted also, except for special cases where its value would be negative or not defined, the other way around: the electric energy saved if we accept a one point performance drop.

3.3. Results

Quantitative results for SLU tasks are evaluated with the Concept Error Rate (CER)⁶ for PortMEDIA and MEDIA tasks, and with accuracy for FSC. We add in tables the number of parameters, and the metrics for evaluating the computational cost of our models, introduced in the previous section.

3.3.1. Preliminary experiments for French

For reaching experimental conditions where we have resources available in advance for our French target task (MEDIA), we trained SLU models for the PortMEDIA task. These models will be then used for pre-initializing models for MEDIA. Results on PortMEDIA are reported in table 1. We trained models with two basic features: spectrograms (indicated with **spectro**) and with features generated with the SSL model *w2v2-fr 7k* from *LeBenchmark*. As we can see, the best results are obtained always with the most expensive strategy *3 steps*. However, using as input *w2v2-fr* features, the difference between *3 steps* and *1 step* models is less than 1 CER point, while the latter strategy is quite cheaper both in terms of training time and energy consumption: 54.43% *kWh* cost reduction with only 4.31% performance loss. Even better CER results can be obtained, and at a lower final cost (69.52% *kWh* cost reduction with 1.69% performance gain),

⁴<https://codecarbon.io>

⁵Available on the European Environment Agency website.

⁶Computed aligning with the edit distance the gold and the predicted concept sequences.

using features computed with the *w2v2-fr* SSL model fine-tuned on the SLU MEDIA task. Supervised fine-tuning is performed like in [23], the SSL model training is resumed using MEDIA data for training and minimizing the CTC loss with respect to the SLU output (see the appendix A.2.3 in [23]). These results are shown in the last line of table 1 (*w2v2-fr slu* features). These results are marked with *1 step + l* in the column **Train stg** to keep into account the fine-tuning of the SSL model, which requires 100 hours of training on 4 GPUs (+100h *x4 GPU* in the tables) of a 318M parameters model. Since this model is never a component of the final SLU model in our experiments, it is indicated in the “# params” column in parenthesis. Based on estimations in [23], we can estimate the cost of our fine-tuning to 92.720 *kWh*, or 4729 *gCO₂*, these values are reported in table headers. The fine-tuning cost dominates the cost for training the SLU model. As we mentioned above however, this fine-tuning is performed only once, on the MEDIA task, and with the purpose of sharing such models as additional resources, like models of *LeBenchmark*. We made this choice both to save resources overall, and because PortMEDIA SLU models are used to pre-initialize SLU models trained for MEDIA as a transfer learning strategy. As we anticipated in previous sections thus, using available SSL models it is possible to obtain competitive performance on the SLU task at a lower cost (*3 steps* vs. *1 step* strategies). Logically, results are even better if a fine-tuned SSL model is available. Thanks to this first set of experiments, we have SLU models available, in addition to SSL models, to train SLU models at a lower cost on our French target task, MEDIA. In a real application scenario, it would be desirable that such resources, including fine-tuned models, were available in advance, and they are exploited by using the same SLU system for training target SLU models, which is what we do in this work.

3.3.2. Results on English FSC

Results on the FSC corpus are given in table 2. As we can see they follow the same trend as results on PortMEDIA. All models obtain very good results on this task, even if we did not perform any fine tuning of parameters on this task, the system is used as in [23], which is tuned on MEDIA. Beyond that, using features from the *XLSR53* model, we obtain very competitive results even with the model trained full end-to-end (*1 step*), with a 37.45% *kWh* cost reduction with respect to the *3 steps* strategy and just 0.77% performance loss. Results are even better, and obtained at a lower final cost (3h23’ for training, 65.82% *kWh* cost reduction and 0.05% performance gain), using a model trained with features from a SSL model tuned on the FSC SLU output. Again, fine-tuning time and cost dominate SLU model cost, but it is intended to be done once and for all, and for producing resources that will be made available for avoiding to repeat such process. Additionally, while we fine-tune a model of 315M parameters on FSC data only, that is 14.7 hours of speech, a state-of-the-art model such as [31] pre-trains ASR and BERT-base models, roughly 287M parameters, on 75k hours of speech, then use such models as components in the final SLU system, which is also fine-tuned on the FSC data.

3.3.3. Results on French MEDIA

Results on the MEDIA corpus are reported in table 3. In the block **Base features** we show results obtained with the same experimental conditions used for PortMEDIA. These results confirm that a competitive SLU model can be obtained at lower training cost: the *1 step* strategy, compared to the *3 steps* strategy, with *w2v2-fr* features, gives a 40.23% *kWh* cost reduction and a

Corpus: MEDIA, Metric: Concept Error Rate (CER ↓)							
Train stg	# params (M)	Input	kWh (gCO ₂)	kWh/p	Train T	DEV	TEST
Base Features							
[23] 3 steps	10.15	spectro	-	-	57h	29.07	31.10
3 steps	10.15	spectro	6,651 (314)	0,273	56h55'	28.35	28.95
2 steps	10.15	spectro	4,417 (225)	0,173	40h52'	32.04	32.85
1 step	10.15	spectro	2,407 (123)	\mathcal{M}_2	22h16'	46.57	44.50
[23] 3 steps	12.16	w2v2-fr	-	-	36h	17.25	16.25
3 steps	12.16	w2v2-fr	3,597 (183)	0,550	36h01'	18.69	16.14
2 steps	12.16	w2v2-fr	2,445 (125)	0,116	24h29'	18.24	16.23
1 step	12.16	w2v2-fr	2,150 (110)	\mathcal{M}_2	21h32'	19.68	18.77
Fine-tuned features (+100h x4 GPU ⇒ 92.720 kWh / 4729 gCO₂)							
2 steps +1	12.16 (+318)	w2v2-fr slu	2,569 (131)	∞	27h28'	14.25	13.78
1 step +1	12.16 (+318)	w2v2-fr slu	2,529 (129)	∞	27h02'	14.16	13.26
[23] ^(*) 3 steps +1	12.16 (+318)	w2v2-fr asr	-	-	36h	14.58	13.78
Transfer learning							
1 step +PM	12.16	w2v2-fr	2,420 (123)	0,125	25h04'	18.27	16.61
Transfer learning + fine-tuned features (+100h x4 GPU ⇒ 92.720 kWh / 4729 gCO₂)							
1 step +1 +PM	12.16 (+318)	w2v2-fr slu	2,026 (103)	\mathcal{M}_2	19h23'	13.59	13.21
State of the art							
[11]	(+)	MFCC	-	-	-	16.1	13.6
[5]	318	w2v2-fr slu ^(**)	-	-	-	-	11.2

Table 3: Results on the French corpus MEDIA. (*): fine tuning for these results was performed on ASR output. (**): features fine-tuned in multiple stages (ASR, SLU) on the MEDIA corpus. (+): we could not estimate as the layer size is not reported.

16.29% performance loss. In the block **Fine-tuned features** we show results obtained with features from the w2v2-fr SSL model fine-tuned on the SLU output of the MEDIA task (w2v2-fr slu). It is interesting to note that the model trained full end-to-end (1 step +1) obtains better results than the model trained in 2 steps. This is due to the fact that the full end-to-end model can be trained with a more aggressive training strategy, in particular a lower dropout regularization. These settings are not effective for the 2-steps model, intuitively because they “erase” the information provided by the pre-trained encoder, bringing the model far from the optimum. Since the SSL model is fine-tuned on the same task as the final SLU model, it is not surprising that the SLU model obtains very competitive results at a lower cost: 29.69% kWh cost reduction and 17.84% performance gain with respect to the 3 steps strategy. Indeed, compared to the last state-of-the-art models on MEDIA (**State of the art** in the table), our results are better than [11], which is the best end-to-end model so far, and not far from [5] which used a cascade system, but our results are obtained at a lower training cost compared to the literature. [11] and [5] do not mention the computational cost of their model, however from what is reported in their work, we estimate a higher computational cost than the one needed overall for our models. [5] in particular uses the wav2vec 318M parameter model as a component of the SLU system and performs several fine-tuning processes like the one we performed once.

Results reported so far prove that SSL models allow for reaching very competitive performance even with a full end-to-end training strategy (1 step), thus for the following experiments we use only this cheaper strategy.

In the blocks **Transfer learning** and **Transfer learning + fine-tuned features** of table 3, we show respectively results obtained with transfer learning from the PortMEDIA task, and results obtained with transfer learning from the same task using features fine-tuned on the SLU MEDIA task.

With transfer learning alone, using a SLU model trained on PortMEDIA as starting point for training a model on MEDIA (+PM), we obtain a substantial improvement on the test set: 18.77 vs. 16.61, and only 2.91% performance loss with respect to the 3 steps strategy (16.14% CER) with 32.72% kWh cost reduction. This could be obtained without any additional cost assuming a situation where SLU models for PortMEDIA are

available in advance. The value kWh/p for models using w2v2-fr features as input is computed with respect to this model (the cheapest one with this input).

With transfer learning and using fine-tuned features (**Transfer learning + fine-tuned features** block), while there is an improvement on the development set (14.16 vs. 13.59), there is basically no gain on the test set (13.26 vs. 13.21). We think this is due to the fact that, since the SSL model is already fine-tuned on the SLU task, the small contribution of PortMEDIA through the transfer learning (PortMEDIA is even smaller than MEDIA) does not add much information that is not already contained in the fine-tuned features. This model has still the advantage of being cheaper: 19h23' vs. 27h02' for training, that is 19.88% kWh cost reduction with respect to the 1 step +1 model trained with fine-tuned features but without transfer learning. The same model allows for a 43.67% cost reduction and 18.15% performance gain with respect to the 3 steps strategy, again assuming conditions where SLU models for PortMEDIA and fine-tuned models are available in advance. The value kWh/p for models using w2v2-fr slu features is computed with respect to this model.

As a final remark, we would like to point out that we were not led “down the garden path” blindly listening stochastic parrots [25], our SLU models proved to be state-of-the-art with basic input features [10], and they were improved in [23], while being generic enough to tackle any sequence prediction problem from speech.

4. Conclusions

In this paper we analysed simple training strategies for SLU models, with the aim of reducing the computational cost of the model training phase, while keeping competitive performance. Our results show that, through the use of pre-trained SSL models, it is possible to attain these objectives. Adding the computational cost for fine-tuning an SSL model on the target SLU task, we obtain the second best result of the literature on MEDIA, and we equal the state-of-the-art on the FSC task, both with a full end-to-end model trained in one shot. While fine-tuning is relatively costly, overall our models are less expensive than the best state-of-the-art models. In order to have

a comprehensive understanding of how results on a given task are obtained, it would be desirable that the community adopt as a standard practice the evaluation of models also in terms of resource consumption, especially as consequence of the more and more frequent use of huge and expensive models.

5. Acknowledgements

This work has been realized using the HPC computational resources of GENSI - IDRIS, contract number AD011011615R1. This work was partially supported by the JCJC CREMA project (*Coreference RESolution into MACHine translation*), funded by the French Agency for Research (ANR), contract number ANR-21-CE23-0021-01.

6. References

- [1] R. De Mori, *Spoken Dialogues with Computers*. Orlando, FL, USA: Academic Press, Inc., 1997.
- [2] C. Raymond, F. Béchet, R. De Mori, and G. Damnati, "On the use of finite state transducers for semantic interpretation," *Speech Communication*, vol. 48, no. 3-4, pp. 288–304, 4 2006.
- [3] S. Hahn, M. Dinarelli, and et al., "Comparing stochastic approaches to spoken language understanding in multiple languages," *IEEE TASLP*, vol. 99, 2010.
- [4] A. Caubrière, S. Ghannay, and et al., "Error analysis applied to end-to end spoken language understanding," in *In ICASSP*, Barcelona, Spain, May 2020.
- [5] S. Ghannay, A. Caubrière, and et al., "Where are we in semantic concept extraction for Spoken Language Understanding? ☆," in *In SPECOM 2021*, Saint Petersburg, Russia, Sep. 2021.
- [6] D. Serdyuk, Y. Wang, and et al., "Towards end-to-end spoken language understanding," *CoRR*, vol. abs/1802.08395, 2018.
- [7] T. Desot, F. Portet, and M. Vacher, "SLU FOR VOICE COMMAND IN SMART HOME: COMPARISON OF PIPELINE AND END-TO-END APPROACHES," in *ASRU Workshop*, Sentosa, Singapore, Singapore, Dec. 2019.
- [8] L. Lugosch, M. Ravanelli, and et al., "Speech model pre-training for end-to-end spoken language understanding," 2019.
- [9] A. Caubrière, N. A. Tomashenko, and et al., "Curriculum-based transfer learning for an effective end-to-end spoken language understanding and domain portability," *CoRR*, vol. abs/1906.07601, 2019.
- [10] M. Dinarelli, N. Kapoor, B. Jabaian, and L. Besacier, "A data efficient end-to-end spoken language understanding architecture," 2020.
- [11] V. Pelloin, N. Camelin, and et al., "End2End Acoustic to Semantic Transduction," in *In ICASSP*, Toronto, ON, Canada, Jun. 2021.
- [12] T. Desot, F. Portet, and M. Vacher, "End-to-end spoken language understanding: Performance analyses of a voice command task in a low resource setting," *Computer Speech & Language*, vol. 75, 2022.
- [13] Y. Qian, R. Ubale, and et al., "Exploring asr-free end-to-end modeling to improve spoken language understanding in a cloud-based dialog system," in *2017 IEEE ASRU Workshop*, 2017.
- [14] E. Palogiannidi, I. Gkinis, G. Mastrapas, P. Mizera, and T. Stafylakis, "End-to-end architectures for asr-free spoken language understanding," 2020.
- [15] D. Serdyuk, Y. Wang, and et al., "Towards end-to-end spoken language understanding," 2018.
- [16] M. Radfar, A. Mouchtaris, and S. Kunzmann, "End-to-end neural transformer based spoken language understanding," 2020.
- [17] S. Schneider, A. Baevski, and et al., "wav2vec: Unsupervised pre-training for speech recognition," *CoRR*, vol. abs/1904.05862, 2019.
- [18] A. Baevski, Y. Zhou, and et al., "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *NeurIPS*, vol. 33, 2020.
- [19] W. Hsu, B. Bolte, and et al., "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *CoRR*, vol. abs/2106.07447, 2021.
- [20] C. Lai, Y. Chuang, and et al., "Semi-supervised spoken language understanding via self-supervised speech and language model pretraining," *CoRR*, vol. abs/2010.13826, 2020.
- [21] S. Yang, P. Chi, and et al., "SUPERB: speech processing universal performance benchmark," *CoRR*, vol. abs/2105.01051, 2021.
- [22] S. Evain, H. Nguyen, and et al., "LeBenchmark: A Reproducible Framework for Assessing Self-Supervised Representation Learning from Speech," in *In Interspeech*, Brno, Czech Republic, Aug. 2021.
- [23] S. Evain, M. H. Nguyen, and et al., "Task Agnostic and Task Specific Self-Supervised Learning from Speech with LeBenchmark," in *In NeurIPS*, ser. NeurIPS 2021 Datasets and Benchmarks Track, on-line, United States, Dec. 2021.
- [24] E. Strubell, A. Ganesh, and A. McCallum, "Energy and policy considerations for deep learning in NLP," in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics, Jul. 2019.
- [25] E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell, "On the dangers of stochastic parrots: Can language models be too big?" in *ACM Conference on Fairness, Accountability, and Transparency*. New York, USA: Association for Computing Machinery, 2021.
- [26] T. Parcollet and M. Ravanelli, "The Energy and Carbon Footprint of Training End-to-End Speech Recognizers," Apr. 2021, working paper or preprint.
- [27] A. Conneau et al., "Unsupervised cross-lingual representation learning for speech recognition," *arXiv:2006.13979*, 2020.
- [28] H. Bonneau-Maynard, C. Ayache, and et al., "Results of the French evalda-media evaluation campaign for literal understanding," in *In LREC*. Genoa, Italy: European Language Resources Association (ELRA), May 2006.
- [29] F. Lefèvre, D. Mostefa, and et al., "Leveraging study of robustness and portability of spoken language understanding systems across languages and domains: the PORTMEDIA corpora," in *LREC*, 2012, pp. 1436–1442.
- [30] A. Caubrière, N. Tomashenko, and et al., "Curriculum-based transfer learning for an effective end-to-end spoken language understanding and domain portability," in *Interspeech*, Graz, Austria, Sep. 2019.
- [31] Y. Qian, X. Bian, and et al., "Speech-language pre-training for end-to-end spoken language understanding," *CoRR*, vol. abs/2102.06283, 2021.
- [32] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, Nov. 1997.
- [33] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," in *In ICLR*, Y. Bengio and Y. LeCun, Eds., 2015.
- [34] W. Chan, N. Jaitly, Q. V. Le, and O. Vinyals, "Listen, attend and spell," *CoRR*, vol. abs/1508.01211, 2015.
- [35] A. Vaswani, N. Shazeer, and et al., "Attention is all you need," in *In NIPS*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017.
- [36] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks," in *Proceedings of ICML*. ACM, 2006, pp. 369–376.

- [37] M. Dinarelli, A. Moschitti, and G. Riccardi, "Re-ranking models based on small training data for spoken language understanding," in *Conference of Empirical Methods for Natural Language Processing*, Singapore, Singapore, aug 2009, pp. 1076–1085.
- [38] S. Quarteroni, G. Riccardi, and M. Dinarelli, "What's in an ontology for spoken language understanding," in *Interspeech*, Brighton, U.K., 2009.
- [39] M. Dinarelli, A. Moschitti, and G. Riccardi, "Concept segmentation and labeling for conversational speech," in *Interspeech*, Brighton, U.K., 2009.
- [40] M. Dinarelli, "Spoken language understanding: from spoken utterances to semantic structures," Ph.D. dissertation, International Doctoral School in Information and Communication Technology, Dipartimento di Ingegneria e Scienza dell' Informazione, via Sommarive 14, 38100 Povo di Trento (TN), Italy, March 2010.
- [41] M. Dinarelli and I. Tellier, "Improving recurrent neural networks for sequence labelling," 2016.
- [42] Y. Dupont, M. Dinarelli, and I. Tellier, "Label-dependencies aware recurrent neural networks," in *Proceedings of the 18th International Conference on Computational Linguistics and Intelligent Text Processing*. Budapest, Hungary: Lecture Notes in Computer Science (Springer), April 2017.
- [43] M. Dinarelli, V. Vukotic, and C. Raymond, "Label-dependency coding in Simple Recurrent Networks for Spoken Language Understanding," in *Interspeech*, Stockholm, Sweden, Aug. 2017.