



HAL
open science

Mediapi-RGB: Enabling Technological Breakthroughs in French Sign Language (LSF) Research through an Extensive Video-Text Corpus

Yanis Ouakrim, Hannah Bull, Michèle Gouiffès, Denis Beautemps, Thomas Hueber, Annelies Braffort

► To cite this version:

Yanis Ouakrim, Hannah Bull, Michèle Gouiffès, Denis Beautemps, Thomas Hueber, et al.. Mediapi-RGB: Enabling Technological Breakthroughs in French Sign Language (LSF) Research through an Extensive Video-Text Corpus. VISAPP 2024 - 19th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications, Feb 2024, Rome, Italy. 10.5220/0012372600003660 . hal-04494094

HAL Id: hal-04494094

<https://hal.science/hal-04494094>

Submitted on 11 Mar 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.

Mediapi-RGB: Enabling Technological Breakthroughs in French Sign Language (LSF) Research through an Extensive Video-Text Corpus

Yanis Ouakrim^{‡,2} Hannah Bull^{*1} Michèle Gouiffès¹
Denis Beautemps² Thomas Hueber² Annelies Braffort¹

¹LISN, Univ. Paris-Saclay, CNRS, 91405 Orsay, France

²Univ. Grenoble Alpes, GIPSA-Lab, CNRS, F-38000 Grenoble, France

{yanis.ouakrim, hannah.bull, annelies.braffort, michele.gouiffes}@lisn.upsaclay.fr,
{yanis.ouakrim, denis.beautemps, thomas.hueber}@gipsa-lab.grenoble-inp.fr

Abstract

We introduce Mediapi-RGB, a new dataset of French Sign Language (LSF) along with the first LSF-to-French machine translation model. With 86 hours of video, it is the largest LSF corpora with translation. The corpus consists of original content in French Sign Language produced by deaf journalists, and has subtitles in written French aligned to the signing. The current release of Mediapi-RGB is available at the Ortolang corpus repository¹, and can be used for academic research purposes. The test and validation sets contain 13 and 7 hours of video respectively. The training set contains 66 hours of video that will be released progressively until December 2024. Additionally, the current release contains skeleton keypoints, sign temporal segmentation, spatio-temporal features and subtitles for all the videos in the train, validation and test sets, as well as a suggested vocabulary of nouns for evaluation purposes. In addition, we present the results obtained on this corpus with the first LSF-to-French translation baseline to give an overview of the possibilities offered by this corpus of unprecedented caliber for LSF. Finally, we suggest potential technological and linguistic applications for this new video-text dataset.

*Equal contribution

¹<https://www.ortolang.fr/workspaces/mediapi-rgb>

1 Introduction

In comparison to written and spoken corpora, there is a lack of sign language (SL) corpora, particularly large corpora with native or near-native signers in natural settings. Unfortunately, the existence of such large annotated corpora is a prerequisite for the automatic processing of sign languages. This is especially true for translation (0) or sentence alignment (0), where deep learning models need to build an understanding of both the source sign language and the target vocal language.

To this day, there are very few sign language corpora translated into vocal languages. And when they do exist, they are often very small (0) due to the fact that they are often produced in laboratories, such as (0; 0), and are therefore expensive to acquire, translate and annotate. To face this challenge, many researchers rely on interpreted sign language data, i.e. a reverse translation from a vocal language to a sign language., e.g. Phoenix-14-T (0), SWISSTXT-WEATHER and SWISSTXT-NEWS (0), Content4all (0) and BOBSL (0) datasets. These corpora are derived from television programs whose vocal language has been transcribed into subtitles and also interpreted into a given sign language. Researchers using these corpora are working on the assumption that the sign language interpretation retains enough information so that it is possible for a translation model to recover the subtitle text from the interpreted sign language. Moreover, there are differences between interpreted and non-interpreted language (0) due to source language interference and time constraints. SL

content from real-time interpretation tends to closely follow the grammatical structure of the spoken language due to strong time constraints (0). There is also some evidence of differences between hearing and deaf interpreters (0), and more generally differences between hearing non-native signers, deaf non-native signers and deaf signers (0). Therefore, the lack of representation of native or near-native deaf signers outside of laboratory conditions is also a key issue.

Datasets of non-interpreted sign language along with translation are very rare; for a recent SL translation challenge, (0) release a dataset with 19 hours of original Swiss-German SL content. Mediapi-Skel, a first French SL (LSF) dataset of 27 hours, has been released with 2D keypoints data aligned with subtitles (0). Mediapi-RGB aims to complement existing corpora by providing 86 hours of video of SL content outside of laboratory conditions with a high representation of deaf native signers.

Existing translated French sign language corpora are of very limited duration (0). Mediapi-RGB is therefore the largest translated LSF corpus, enabling to train machine translation models. Along with the presentation of the corpus we apply a simple transformer based machine translation baseline to give an overview of the translation possibilities with this new corpus following existing works on German Sign Language (DGS) (0), British Sign Language (BSL) (0) and American Sign Language (ASL) (0). The two main contributions of this paper are the following:

- The provision of the largest translated LSF corpus to date and one of the largest compared to other sign language corpora around the world.
- The first machine translation model from LSF video to French text.

The remainder of the paper is organized as follows. Section 2 provides an overview of the dataset and section 3 gives information about its production. Then, section 4 proposes a translation baseline and its results on Mediapi-RGB. Opportunities and limitations of this new dataset are discussed in Section 5.

2 Dataset overview

Mediapi-RGB is a large corpus (86 hours) of original LSF content available for academic research purposes. This corpus can be used for numerous research applications, including training or evaluating SL retrieval, recognition or translation models.



Figure 1: **Mediapi-RGB source data.** The source data consists of journalistic content in LSF, along with translated French subtitles aligned to the signed content. (Subtitles are not embedded in the video pixels).

We provide a summary of the Mediapi-RGB dataset and compare this corpus with the Mediapi-Skel dataset.

2.1 Dataset Content and Statistics

The source of the data is the French media association *Média’Pi!*,² the same source of the data used to create Mediapi-Skel (0). *Média’Pi!*, is a generalist and independent online media launched in April 2018. It offers articles and reports on national and international news as well as topics related to the deaf community in LSF and French. It operates on a subscription basis. The quantity and quality of original content produced by *Média’Pi!* (noted episodes or original videos in the paper) is difficult to find outside of laboratory-produced corpora, and the diversity of scenarios better reflects the diversity of real-world SL videos. The LSF produced by the deaf journalists follows a natural structure. In a second stage, the LSF content is translated into written French and aligned to the LSF video in the form of subtitles. The alignment between written subtitles and LSF video is accurate. This is in contrast to interpreted data, where there is often a lag between the written subtitles (aligned to the audio) and the sign language interpretation. Figure 1 shows a screenshot of a video from the *Média’Pi!* website along with its associated subtitle in French (“Last Saturday, 19,000 doses of vaccine were released to vaccinate the entire major population of the area”).

The dataset contains 1,230 videos, representing a total of 86 hours of LSF produced between 2017 and 2022. In order to protect the economic model of

²<https://media-pi.fr/>

Média’Pi!, we are only able to publish videos of information and sports programmes with a time delay of three years, and these videos may only be used for academic research purposes. We have decided to use videos dating from 2017-2018 in the validation set, and videos from 2019 in the test set, so that models trained on other corpora can already be evaluated on Mediapi-RGB. Other videos from 2020-2022 will be progressively released in the training set until December 2024. We are nevertheless able to release keypoints data, video features, automatic sign segmentation and subtitles of the training set videos. We also propose a 4k vocabulary of common and proper nouns in the Mediapi-RGB subtitles for evaluation purposes. Table 1 gives information on the split between train, validation and test sets.

	Train	Val	Test	Total
Fully released	No	Yes	Yes	-
# videos	950	74	206	1230
# subtitles	37651	4373	8060	50084
# hours (full video)	66.1	7.2	12.5	85.9
# hours (subtitles)	52.3	4.9	10.8	68.0

Table 1: **Train, validation and test split.** The RGB training videos will be released progressively until Dec. 2024. Features and skeleton keypoints on the training set are available.

2.2 Relationship to Mediapi-Skel

The Mediapi-RGB dataset is intended as a replacement rather than a complementary dataset to Mediapi-Skel. This is because the training set of Mediapi-Skel partially overlaps with the test partition of Mediapi-RGB. The current release of the training set of Mediapi-RGB contains skeleton keypoints for the train, validation and test sets, and so skeleton-based models can currently be trained and evaluated on Mediapi-RGB rather than Mediapi-Skel. The main advantages of Mediapi-RGB are that: 1) it’s a bigger corpus (86 hours vs. 27 hours), 2) the original RGB videos are available on the validation and test sets, and will soon be released for the training set, and 3) additional data such as Mediapipe keypoints (0), Video Swin Transformer features (0) and automatic sign temporal segmentations are available. The main disadvantage of Mediapi-RGB in comparison to Mediapi-Skel is the fact that there are fewer signers,

due to the exclusion of certain types of programmes including interviews. A comparison of the statistics of Mediapi-Skel and Mediapi-RGB can be found in Table 2.

	Mediapi-Skel	Mediapi-RGB
Global statistics		
# subtitled episodes	368	1230
# hours	27	86
# frames	2.5 million	7.7 million
Video statistics		
Resolution (# vids.)	1080p (327) 720p (41)	2160p (28) 1080p (1182) 720p (18) 480p (2)
Framerate (# vids.)	30 fps (111) 25 fps (242) 24 fps (15)	25 fps
Avg. length of vid.	4.5 minutes	4.2 minutes
# signers	> 100	> 10
Text statistics		
# subs.	20 187	50 084
Avg. length of sub.	4.2 seconds 10.9 words	4.9 seconds 12.2 words
Vocab. (tokens)	17 428	35 599
Vocab. (nouns+verbs+adj.)	14 383	27 343

Table 2: Descriptive statistics of our datasets Mediapi-Skel and Mediapi-RGB.

3 Dataset collection

This section describes the dataset collection, as well as various post-processing steps to facilitate usage of Mediapi-RGB. We temporally crop the videos using the start and end times of each subtitle (Sec. 3.1) and spatially crop the signers (Sec. 3.2). This creates video-text sentence-like pairs, with the signer centered within a square crop of 444×444 pixels at 25fps. We remove duplicates to ensure that there are no identical videos across the train, validation and test splits (Sec. 3.3). Although we can not currently release the training set videos, we can release the corresponding processed data, including OpenPose (0) and Mediapipe Holistic (0) keypoints (Sec. 3.4), Video Swin Transformer features (Sec. 3.5), sign segmentation (Sec. 3.6), original




Figure 2: **Duration of subtitles.** The distribution of the durations of the extracted subtitle clips (without 0.5s padding).




Figure 3: **Duration of episodes.** The distribution of the durations of *Média'Pi!* original episodes (without any modification).

and processed subtitles (Sec. 3.7), as well as a proposed noun vocabulary for evaluation purposes (Sec. 3.8).

3.1 Temporally Cropping Subtitles

In Mediapi-*RGB*, the subtitles are aligned to the signing. Almost all of the *Média'Pi!* episodes are produced in LSF, then subsequently subtitled in written French. There are rare cases where a hearing person is interviewed in spoken language and this is interpreted into LSF. In these cases, the subtitles correspond to the original audio, but are aligned to the signing, and the audio track is removed. We temporally crop the videos using the subtitle timestamps, adding 0.5s on each side as padding. This creates sentence-like video-text pairs. The distribution of subtitle durations is shown in Fig. 2, and the distribution of the duration of the original episodes is shown in Fig. 3.

3.2 Spatially Cropping Signers

As the *Média'Pi!* videos capture signers in different positions and orientations, we automatically extract bounding boxes around the most likely signer in each of the extracted subtitle crops (Sec. 3.1). To do this, we follow the methodology described in (0) with available online code.³ We extract the 2D OpenPose (0) keypoints of each person in the videos, omit the legs and feet keypoints, track each person between consecutive frames, impute missing skeleton keypoints using past

³https://github.com/hannahbull/clean_op_data_sl

or future frames, temporally smooth keypoints using a Savitzky-Golay filter and omit unlikely signers such as people with occluded hands, people with hands that hardly move or people in the background. In the case of multiple potential signers, we then choose the most likely signer based on a metric computed by multiplying the hand size and the variation of wrist movement of the dominant hand. We then use a static square crop around this most likely signer for the duration of each subtitle. This square crop is then resized to 444×444 pixels.

3.3 Removing Duplicate Videos

In the initial dataset, there were a large amount of duplicate videos, generally some extracts from an original video (for advertisement or quoting) or recurrent opening/closing titles. These duplications have to be removed to avoid biases. To that purpose, we select the videos that share the same subtitle and measure the spatio-visual similarity between them using the Video Swin features, described further in Sec. 3.5. Letting $\{\mathbf{a}_1, \dots, \mathbf{a}_n\}$ be the sequence of visual feature vectors of a video A and $\{\mathbf{b}_1, \dots, \mathbf{b}_m\}$ be the sequence of visual feature vectors of a video B , we compute:

$$S_{A,B} = \max_{i \in 1, \dots, n, j \in 1, \dots, m} \frac{\mathbf{a}_i \cdot \mathbf{b}_j}{\|\mathbf{a}_i\| \|\mathbf{b}_j\|}. \quad (1)$$

If the maximum cosine similarity $S_{A,B}$ in (1) is above a threshold $T = 0.95$, the two videos are considered to be duplicates. Once duplicated crops A or B are detected, the longer one is kept since it is more




Figure 4: **Illustration of OpenPose and Mediapipe Holistic keypoints.** Illustration of the face, body and hand keypoints detection using both OpenPose (0) (left) and Mediapipe Holistic (0) (right).

likely to correspond to the original video. Crops with a similarity between $T = 0.85$ and $T = 0.95$ are considered as *potential* duplicates and have therefore been grouped in the same split to avoid biases. Thresholds were set on the basis of manual observations.

3.4 Skeleton Keypoints Extraction

Skeleton keypoints are useful inputs for many automatic SL processing tasks, such as hand or face cropping (0; 0), generating SL (0; 0), or as inputs to improve recognition methods (0; 0). Thus, OpenPose (0) and Mediapipe Holistic (0) keypoints are provided for the body, the hands and the face. Fig. 4 shows an example of the face, body and hand keypoints detection with both Mediapipe Holistic and OpenPose.

3.5 Spatio-temporal Features Extraction

Features are a useful input for automatic SL processing systems, because they allow models to be trained more rapidly than end-to-end systems using RGB video frames as input. For example, features are used directly to train an automatic alignment model in (0) and to search for lexical signs in (0).

To extract features, we use the Video Swin transformer model (0) trained for the task of sign recognition using dense automatic annotations acquired using the methods described in (0). Superior performance of this model over I3D models is shown in (0) for sign classification. The input to the Video Swin transformer model is a temporal context window of 16 frames, and the output is a vector of features of dimension 768. We




Figure 5: **Sign temporal segmentation.** Example of sign segmentation on an extract of video from Mediapipe-RGB. The blue line denotes the predicted output scores and the green blocks denote a binary threshold for scores above 0.5.

extract these features at stride 1. Due to 0.5s padding applied on each of the subtitle timestamps during temporal cropping, all of the extracted videos contain at least 16 frames.

3.6 Sign Temporal Segmentation

In (0), the authors train a model to automatically segment signs using annotated sign glosses from BSL Corpus (0). This model inputs I3D features from (0) and outputs change-points between signs. The model tends to recognise changes in handshape, and thus over-segments fingerspelling. We use available online code⁴ to segment signs in the Mediapipe-RGB subtitle crops. Although trained on BSL data, this model provide valuable approximations of sign boundaries in LSF. Fig. 5 shows an illustration of the sign segmentation model on a Mediapipe-RGB video. The sign segmentation model recognises transitions between signs due to visual cues such as changes in hand shape.

3.7 Text Processing

We provide the raw subtitle texts for all of the videos in the train, validation and test sets. Additionally, we extract the part-of-speech of each subtitle word⁵. The total vocabulary size is 35,599 and the vocabulary size of nouns, verbs and adjectives is 27,343 (Tab. 2). Other parts of speech such as prepositions, determinants and

⁴<https://github.com/RenzKa/sign-segmentation>

⁵<https://huggingface.co/gilf/french-postag-model>




Figure 6: **Top 20 most frequent words in noun vocabulary.** The frequency of the top 20 noun occurrences.

adverbs do not have associated lexical signs. Linguistic elements of this type are carried by non-manual components or movement features.

3.8 Noun Vocabulary

We propose a vocabulary of 3,894 common and proper nouns for evaluation purposes. This vocabulary corresponds to a list of nouns from the Mediapi-RGB subtitles, appearing at least 5 times. These nouns were lemmatized using the spaCy lemmatizer trained on French news text⁶. Fig. 6 plots the frequency of the top 20 most commonly used nouns in this corpus. Fig. 7 shows the distribution of the number of words in each subtitle, as well as the number of words in the vocabulary in each subtitle.

4 Translation baseline

Following work conducted on other sign languages (0; 0), we apply a transformer (0) based translation baseline to Mediapi-RGB for LSF-to-French translation. An overview of the architecture of this baseline is given in Fig. 8. On the sign language video side, we use the features obtained with a frozen Video Swin transformer model (0) trained on British Sign Language (BSL) presented in section 3.5. On the text side, we use a tokenizer adapted to French⁷. Finally, the model is case-sensitive.

⁶https://spacy.io/models/fr#fr_core_news_md

⁷CamemBERTTokenizer made available by Hugging face based on SentencePiece (0)

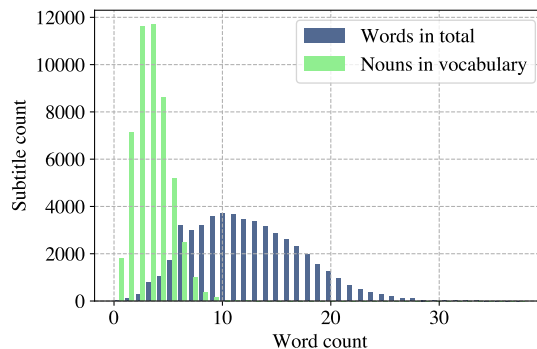


Figure 7: **Number of words in subtitles.** In blue, the distribution of the total number of words in the Mediapi-RGB subtitles; and in green, the distribution of the number of the nouns in our vocabulary in each subtitle (see Sec. 3.8 for details on this vocabulary).

4.1 Implementation Details

Our architecture comprises two transformer encoder layers and one transformer decoder layer. Each transformer block has 8 attention heads. Output of the encoder block is a vector of 768 values. The fully connected layer at the end of the decoder part has a dimension of 64. We train the model with a batch size of 32 and a RMSProp optimizer with a learning rate of 10^{-3} . We use learning rate scheduling: the learning rate is halved with no improvement in validation loss over 7 epochs. Similarly, we use early stopping: training stops after 27 epochs after no improvement in validation loss for 15 epochs. With this mechanism, the result shown in the table corresponds to weights obtained after the 12th epoch. The dropout is set to 0.5. We provide source code for the training and evaluation of this model⁸ as well as the weights of trained models. Future work could focus on hyperparameter tuning and improving this baseline architecture.

4.2 Evaluation and results

Table 3 gives obtained BLEU score of the proposed baseline on the Mediapi-RGB test set. We use the standard SacreBLEU (0) parameters for the BLEU score computation (tokenizer 13a). As BLEU scores depend on language structure, meaningful comparison can only

⁸https://github.com/youakrim/sign_language_translation_Mediapi-RGB

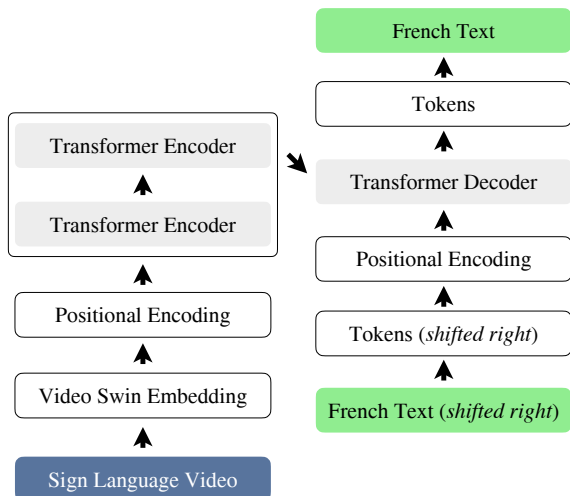


Figure 8: **Translation baseline architecture.** Based on the *Transformer* encoder-decoder architecture presented in (0).

	BLEU-1	BLEU-2	BLEU-3	BLEU-4
Ours	21.56	10.76	6.40	4.14

Table 3: BLEU scores obtained with our translation baseline on the Mediapi-RGB test set.

be achieved with the same target language⁹. To the best of our knowledge, this work is the first to adress LSF-to-French translation, therefore, no comparison to results from other work or method is possible so far. Table 4 provides translation examples with associated ground truth along with English translations obtained with on-line translator DeepL¹⁰. We can note that the model has learnt semantic links: predictions mostly employ words from the same lexical field than the corresponding ground truth.

5 Opportunities and limitations

In this section, we discuss opportunities and limitations of the Mediapi-RGB dataset for technological applications (Sec. 5.1) and for research in SL linguistics

⁹English translations given purely for information purposes in Table 4, show this strong link between BLEU score and target language. For the four given examples, the obtained BLEU score if the target language was English would have been 11.75 instead of 0.0.

¹⁰www.deepl.com

GT	La Première ministre britannique Theresa May a annoncé sa démission vendredi, (British Prime Minister Theresa May announced her resignation on Friday.)
Pred	Le Premier ministre britannique a annoncé vendredi, le Premier ministre britannique, Jean Castex, (The British Prime Minister announced on Friday, the British Prime Minister, Jean Castex.)
GT	Bonjour, nous sommes le mercredi 11 mars 2019. (Hello, today is Wednesday, March 11, 2019.)
Pred	Bonjour! Nous sommes le lundi 11 mars 2019, (Hello! Today is Monday, March 11, 2019.)
GT	Ehpad : intoxication alimentaire mortelle. (French nursing home: fatal food poisoning.)
Pred	L'assurance-nil est un peu plus de retraite. (The insurance-nil is a little more retirement.)
GT	fait de nouveau l'objet d'une enquête, ouverte par le parquet de Paris. (is once again under investigation by the Paris public prosecutor's office.)
Pred	Le procès s'est ouvert le nouveau, le nouveau, en justice de Paris. (The trial opened on the new, new, Paris court.)

Table 4: **Prediction examples with associated ground truth.** Text in brackets corresponds to English translations obtained using online translator DeepL. These examples all obtained a BLEU-4 score of 0.

(Sec. 5.2).

5.1 Applications Perspective

We note that there are historical examples of SL technological applications with little consultation of the deaf community nor practical value (0; 0). Research projects using Mediapi-RGB should take into consideration input from deaf researchers and members of the deaf community in order to ensure practical benefits and to avoid harm. *Média'Pi!* is a project created by deaf journalists to increase access to information in the deaf community. In that regard, we note three potential applications that could take advantage of Mediapi-RGB: anonymization, sign retrieval and automatic translation/subtitling.

One key characteristic of written language is the ability to share information anonymously. It is onerous to present information in SL without also representing the identity of the signer. Solutions such as computer-generated avatars are difficult to implement. MOCAP methods for animating SL avatars are not necessarily anonymous, as individuals can be recognised by their movements and signing style (0). In (0), the authors discuss methods to remove identity cues from SL production. Such methods could be combined with techniques of realistic sign generation (0; 0). Anonymous representations of SL can then be used to communicate factual information such as legal or administrative information, governmental documents or weather reports, or to represent signers who wish to conceal their identity.

Search engines are very efficient for finding written information based on textual queries. Searching for information in SL using text queries or sign queries is very difficult due to the challenges of recognising signs and clustering similar topics in SL videos without written translations. In (0), the authors present a method for searching for SL sequences using free-form textual queries. In (0), the authors search for isolated signs in continuous SL video on BSL Corpus (0) and PHOENIX14-T (0; 0). The Mediapi-RGB corpus can be used to train and evaluate models for retrieval tasks on continuous SL videos using textual or sign queries, as the subtitles of the videos may be used as weak annotation.

Adding subtitles to SL video improves accessibility and comprehension, but is a time-consuming task. Various ways to simplify this task are automatic segmentation into subtitle-units (0), automatic alignment of subtitles to video (0), and eventually automatic translation

of SL video to text. These tasks can be trained and evaluated using the videos in the Mediapi-RGB corpus as demonstrated in section 4.

Users of Mediapi-RGB should be aware of the specificities of this dataset. Models trained on other datasets may not necessarily perform well on Mediapi-RGB, and models trained on Mediapi-RGB may not generalise well to other situations. The *Média'Pi!* videos are of professional quality and are unlikely to be representative of spontaneous conversations in daily life. This can be considered both an advantage and a limitation of Mediapi-RGB. The signers are highly skilled deaf journalists producing examples of eloquent and formal LSF, but models trained on Mediapi-RGB may be less applicable to signers with lower levels of fluency or a more informal register.

5.2 A Sign Linguistics Perspective

One common critique of current large SL corpora for automatic SL processing is the over-representation of interpreted TV data (0). Nevertheless, the differences between interpreted SL data and non-interpreted SL data are not well understood. Comparing Mediapi-RGB with the other available journalistic SL content from interpreted journalistic data sources such as SWISSTXT-NEWS (0), VRT-NEWS (0) and BOBSL (0) may provide some clues on how SL production from deaf journalists differs from interpreted journalistic content from both hearing and deaf interpreters. Increased linguistic knowledge about this differences would help acknowledge and alleviate the biases of models trained on interpreted SL data.

The number of signers is lower in Mediapi-RGB than in Mediapi-Skel, due to the removal of numerous videos involving interviews with members of the public at events. Nevertheless, there are over 10 different signers in Mediapi-RGB, allowing for linguistic comparisons across signers.

6 Conclusion

We presented a summary of the new dataset Mediapi-RGB, a large dataset with 86 hours of content from deaf journalists at the French online media *Média'Pi!*. The videos are in LSF with aligned subtitles in written French. Mediapi-RGB can be used for academic research purpose. We presented a first LSF-to-French

translation baseline, showing that Mediapi-*RGB* is suitable for this research task. This is a stepping stone towards future improvements. We hope that Mediapi-*RGB* can be used to train and evaluate more automatic SL processing tasks, as well as contribute to research in SL linguistics.

Acknowledgements

We thank *Média'Pi* ! for providing the data. This work has been funded by the Bpifrance investment “Structuring Projects for Competitiveness” (PSPC), as part of the Serveur Gestuel project (IVès and 4Dviews Companies, LISN — University Paris-Saclay, and Gipsa-Lab — University Grenoble Alpes).

REFERENCES

- Necati Cihan Camgoz, Simon Hadfield, Oscar Koller, Hermann Ney, and Richard Bowden. Neural sign language translation. In *CVPR*, 2018.
- Hannah Bull, Triantafyllos Afouras, Gül Varol, Samuel Albanie, Liliane Momeni, and Andrew Zisserman. Aligning subtitles in sign language videos. In *ICCV*, 2021.
- Annelies Braffort. Langue des Signes Française : Etat des lieux des ressources linguistiques et des traitements automatiques. In Leonor Becerra, Benoît Favre, Claire Gardent, and Yannick Parmentier, editors, *Journées Jointes des Groupements de Recherche Linguistique Informatique, Formelle et de Terrain (LIFT) et Traitement Automatique des Langues (TAL)*, pages 131–138, Marseille, France, November 2022. CNRS.
- Valentin Belissen, Annelies Braffort, and Michèle Gouiffès. Dicta-Sign-LSF-v2: Remake of a continuous french sign language dialogue corpus and a first baseline for automatic sign language processing. In *Proceedings of the Twelfth International Conference on Language Resources and Evaluation (LREC'20)*, pages 6040–6048, Marseille, France, May 2020. European Language Resource Association (ELRA).
- Jérôme Fink, Benoît Frénay, Laurence Meurant, and Anthony Cleve. Lsfb-cont and lsfb-isol: Two new datasets for vision-based sign language recognition, 2021.
- Jens Forster, Christoph Schmidt, Thomas Hoyoux, Oscar Koller, Uwe Zelle, Justus H Piater, and Hermann Ney. Rwth-phoenix-weather: A large vocabulary sign language recognition and translation corpus. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 3746–3753, Istanbul, Turkey, May 2012. European Language Resource Association (ELRA).
- Necati Cihan Camgoz, Ben Saunders, Guillaume Rochette, Marco Giovanelli, Giacomo Inches, Robin Nachtrab-Ribback, and Richard Bowden. Content4all open research sign language translation datasets. *arXiv preprint arXiv:2105.02351*, 2021.
- Samuel Albanie, Gül Varol, Liliane Momeni, Triantafyllos Afouras, Hannah Bull, Himel Chowdhury, Neil Fox, Rob Cooper, Andrew McParland, Bencie Woll, and Andrew Zisserman. BOBSL: BBC-Oxford British Sign Language Dataset. *arXiv preprint arXiv:2111.03635*, 2021.
- Daria Dayter. Collocations in non-interpreted and simultaneously interpreted english: a corpus study. In *New empirical perspectives on translation and interpreting*, pages 67–91. Routledge, 2019.
- Lorraine Leeson. Making the effort in simultaneous interpreting. In *Topics in Signed Language Interpreting: Theory and Practice*, volume 63, chapter 3, pages 51–68. John Benjamins Publishing, 2005.
- Christopher Stone and Debra Russell. Interpreting in international sign: Decisions of deaf and non-deaf interpreters. 2013.
- Jill P Morford and Martina L Carlson. Sign perception and recognition in non-native signers of asl. *Language learning and development*, 7(2):149–168, 2011.
- Mathias Müller, Sarah Ebling, Eleftherios Avramidis, Alessia Battisti, Michèle Berger, HfH Zurich, Richard Bowden, Annelies Braffort, Necati Cihan Camgöz, Cristina España-Bonet, et al. Findings of the first wmt shared task on sign language translation.
- Hannah Bull, Annelies Braffort, and Michèle Gouiffès. MEDI-API-SKEL - a 2D-skeleton video database of french sign language with aligned french subtitles. In *Proceedings of the Twelfth International Conference on Language Resources and Evaluation (LREC'20)*, pages 6063–6068, Marseille, France, May 2020. European Language Resource Association (ELRA).
- Necati Cihan Camgoz, Simon Hadfield, Oscar Koller, Hermann Ney, and Richard Bowden. Neural sign language translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7784–7793, 2018.
- Laia Tarrés, Gerard I Gállego, Amanda Duarte, Jordi Torres, and Xavier Giró-i Nieto. Sign language translation from instructional videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5624–5634, 2023.
- Camillo Lugaresi, Jiuqiang Tang, Hadon Nash, Chris McClanahan, Esha Uboweja, Michael Hays, Fan Zhang, Chuo-Ling Chang, Ming Guang Yong, Juhyun Lee, Wan-Teh Chang, Wei Hua, Manfred Georg, and Matthias Grundmann. Mediapipe: A framework for

- building perception pipelines. *CoRR*, abs/1906.08172, 2019.
- Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3202–3211, 2022.
- Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. OpenPose: Realtime multi-person 2D pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(1):172–186, 2019.
- Hannah Bull, Michèle Gouiffès, and Annelies Braffort. Automatic segmentation of sign language into subtitle-units. In *ECCVW, Sign Language Recognition, Translation and Production (SLRTP)*, 2020.
- Jie Huang, Wengang Zhou, Qilin Zhang, Houqiang Li, and Weiping Li. Video-based sign language recognition without temporal segmentation. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Bowen Shi, Aurora Martinez Del Rio, Jonathan Keane, Diane Brentari, Greg Shakhnarovich, and Karen Livescu. Fingerspelling recognition in the wild with iterative visual attention. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5400–5409, 2019.
- Lucas Ventura, Amanda Duarte, and Xavier Giró-i Nieto. Can everybody sign now? exploring sign language video generation from 2d poses. *arXiv preprint arXiv:2012.10941*, 2020.
- Ben Saunders, Necati Cihan Camgoz, and Richard Bowden. Everybody sign now: Translating spoken language to photo realistic sign language video. *arXiv preprint arXiv:2011.09846*, 2020.
- Valentin Belissen, Annelies Braffort, and Michèle Gouiffès. Experimenting the automatic recognition of non-conventionalized units in sign language. *Algorithms*, 13(12):310, 2020.
- Tao Jiang, Necati Cihan Camgoz, and Richard Bowden. Skeletor: Skeletal transformers for robust body-pose estimation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- Liliane Momeni, Hannah Bull, KR Prajwal, Samuel Albanie, Gül Varol, and Andrew Zisserman. Automatic dense annotation of large-vocabulary sign language videos. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXV*, pages 671–690. Springer, 2022.
- KR Prajwal, Hannah Bull, Liliane Momeni, Samuel Albanie, Gül Varol, and Andrew Zisserman. Weakly-supervised fingerspelling recognition in british sign language. In *British Machine Vision Conference (BMVC) 2022*, 2022.
- Katrin Renz, Nicolaj C Stache, Neil Fox, Gul Varol, and Samuel Albanie. Sign segmentation with changepoint-modulated pseudo-labelling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3403–3412, 2021.
- Adam Schembri, Jordan Fenlon, Ramas Rentelis, and Keasy Cormier. British Sign Language corpus project: A corpus of digital video data and annotations of British Sign Language 2008–2017 (Third Edition). <http://www.bslcorpusproject.org>, 2017.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Necati Cihan Camgoz, Oscar Koller, Simon Hadfield, and Richard Bowden. Sign language transformers: Joint end-to-end sign language recognition and translation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- Taku Kudo and John Richardson. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. *arXiv preprint arXiv:1808.06226*, 2018.
- Matt Post. A call for clarity in reporting BLEU scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Belgium, Brussels, October 2018. Association for Computational Linguistics.
- Danielle Bragg, Oscar Koller, Mary Bellard, Larwan Berke, Patrick Boudreault, Annelies Braffort, Naomi Caselli, Matt Huenerfauth, Hernisa Kacorri, Tessa Verhoef, et al. Sign language recognition, generation, and translation: An interdisciplinary perspective. In *The 21st International ACM SIGACCESS Conference on Computers and Accessibility*, pages 16–31, 2019.
- Michael Erard. Why sign-language gloves don't help deaf people. *The Atlantic*, <https://www.theatlantic.com/technology/archive/2017/11/why-sign-language-gloves-dont-help-deaf-people/545441/>, 2017.
- Félix Bigand, Elise Prigent, and Annelies Braffort. Retrieving human traits from gesture in sign language: The example of gestural identity. In *Proceedings of the 6th International Conference on Movement and Computing*, pages 1–4, 2019.
- Félix Bigand, Elise Prigent, and Annelies Braffort. Synthesis for the kinematic control of identity in sign language. In *Proceedings of the 7th International Workshop on Sign Language Translation and Avatar Technology: The Junction of the Visual and the Textual: Challenges and Perspectives*, pages 1–6, 2022.

- Amanda Duarte, Samuel Albanie, Xavier Giró-i Nieto, and Gül Varol. Sign language video retrieval with free-form textual queries. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14094–14104, 2022.
- Tao Jiang, Necati Cihan Camgoz, and Richard Bowden. Looking for the signs: Identifying isolated sign instances in continuous video footage. *IEEE International Conference on Automatic Face and Gesture Recognition*, 2021.
- Jens Forster, Christoph Schmidt, Oscar Koller, Martin Bellgardt, and Hermann Ney. Extensions of the sign language recognition and translation corpus rwth-phoenix-weather. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 1911–1916, 2014.