

Digital Object Identifier

Automatic SignWriting Recognition: Combining Machine Learning and Expert Knowledge to Solve a Novel Problem

ANTONIO F. G. SEVILLA^{1,2}, ALBERTO DÍAZ ESTEBAN^{1,3}, AND JOSÉ MARÍA LAHOZ-BENGOECHEA²

¹Software Engineering and Artificial Intelligence, Universidad Complutense de Madrid, c/ Profesor José García Santesmases, 9 28040 Madrid, Spain

²Spanish Linguistics and Literary Theory, Universidad Complutense de Madrid, Facultad de Filología edificio D, c/ Prof. Aranguren s/n, 28040 Madrid, Spain

³Knowledge Engineering Institute, Universidad Complutense de Madrid, Facultad de Psicología, Lateral 2 28223 Pozuelo de Alarcón, Spain

Corresponding author: Antonio F. G. Sevilla (afgs@ucm.es).

The development of the expert system described in this paper is part of the project “Visualizando la SignoEscritura” (Visualizing SignWriting, <https://www.ucm.es/visse>), reference number PR2014_19/01, developed in the Faculty of Computer Science of Universidad Complutense de Madrid, and funded by Indra and Fundación Universia in the IV call for funding aid for research projects with application to the development of accessible technologies. We want to acknowledge the collaboration of the signing community, especially the Spanish Sign Language teachers at Idiomas Complutense and Fundación CNSE.

ABSTRACT Sign languages are viso-gestual languages, using space and movement to convey meaning. To be able to transcribe them, SignWriting uses an iconic system of symbols meaningfully arranged in the page. This two-dimensional system, however, is very different to traditional writing systems, so its automatic processing poses a novel challenge for computational linguistics. In this article, we present a novel problem for the state of the art in artificial intelligence: automatic SignWriting recognition. We examine the problem, model the underlying data domain, and present a first solution in the form of an expert system that exploits the domain knowledge encoded in the data modelization. This system uses an adaptable pipeline of neural networks and deterministic processing, overcoming the challenges posed by the novelty and originality of the problem. Thanks to our data modelization, it improves the accuracy compared to a straight-forward deep learning approach by 17%. All of our data and code are publicly available, and our approach may be useful not only for SignWriting processing but also for other similar graphical data.

INDEX TERMS Computer Vision, Deep Learning, Expert Knowledge, Neural Networks, Sign Language, SignWriting

I. INTRODUCTION

Sign Languages are a family of viso-gestual languages in use by the Deaf and Hard-of-hearing communities. Instead of sound, they rely on hand and body gestures to communicate meaning, making use of the rich possibilities of three-dimensional space and movement to build words and sentences [1], [2]. They are not, as it was sometimes thought by the general population, mere codings of oral language in gestures, but full natural languages with their grammar, vocabulary and evolution [3]–[5]. This recognition as actual human languages has turned them into an object of increasing interest in the research community.

Nonetheless, there is not a standard writing system in use by the signing population. The complex nature of viso-gestual communication, very different to speech, which is based on sound, makes creating or adapting a writing system




FIGURE 1. SignWriting transcription for the sign “History” in Spanish Sign Language, showing the hands near the chin in an initial configuration and moving down and sideways to a finishing configuration.

not a trivial task. A number of proposals exist, most coming from the research community [6]–[8]. A different proposal,

purported to be more user friendly, is SignWriting, which utilizes more of the graphical potential of the blank page to capture signing and its use of space in an expressive and iconic manner [9].

SignWriting uses abstract but recognizable symbols for the hands and other parts of the body, and then places them in 2D space to represent their relative locations in 3D signing space [10]. Arrows and other graphical tricks serve to fully capture the missing third and fourth dimensions (depth and time) in a system that is arguably intuitive for both signers and non-signers, making it not only a valuable recording and communication tool but also very useful for education.

In Fig. 1 an example SignWriting transcription is shown. This sign starts with the hands touching the chin, with the fingers flexed in a “pin” configuration. This configuration, as well as the hand orientation, are represented by the two top hand symbols. The location is signified by the circle, an iconic representation of the head, with a line marking that the chin is the actual point of contact. That there is contact at all is represented by the small asterisk to the left of the head. The hands then move downwards and to the sides, represented by the arrows, and end up fully extended and in a different orientation. In this final position, the hands are filled in black and with the fingers separated from the body of the hand to represent that the palm is facing downwards. This is just a small sample of the richness and complexity of SignWriting, but more can be read in section II or online at <https://www.signwriting.org/>.

The use of graphic properties and spatial relationships makes SignWriting very different from traditional writing systems. Most writing systems in use for oral languages, if not all, are based on the sequential concatenation of characters, with only slight deviations in the co-location of diacritics and sometimes punctuation. In SignWriting, however, “characters” do not occur in a particular one-dimensional order. Instead, they appear in a complex bi-dimensional relationship, where their relative position is not arbitrary but actually represents some spatial meaning. Symbols can be rotated and reflected, and as seen before, filled in with different patterns to represent different orientations. If we count every possible graphical transformation or variation of the characters in SignWriting, there are beyond thirty thousand unique symbols to remember, understand, and be able to produce¹, which again makes it very different from the “usual” writing systems.

These challenges mean that most SignWriting is produced and consumed in graphical format (i.e. images), making it difficult to process with existing language technologies or to combine with oral language resources in equal footing. Existing approaches avoid this problem by using code representations of SignWriting, like SignWriting Markup Language [11], [12] or Formal Signwriting [13], but to be able to computationally process all the SignWriting data that exists

¹Counted in the SignWriting fonts, downloaded from <https://www.signwriting.org/catalog/sw214.html>

in image form, techniques from artificial intelligence and computer vision are required.

We present in this article such a system, capable of understanding instances of SignWriting in image form by extracting a meaningful representation from the raw pixel data. This meaningful representation encodes domain knowledge of SignWriting, and we have developed an intelligent system able to automatically extract it. As far as we can ascertain from the literature, we are the first to process SignWriting images in this way.

For the processing of the graphic data in the SignWriting images, inevitably noisy and very variable, we use deep neural networks, which have very good results in pattern matching and recognition. First, we present a single state-of-the-art neural network to solve the full task, as an example one-shot approach and as baseline for comparison. Improving on the issues that this first solution presents, our proposed system further utilizes the domain knowledge encoded in our data annotation. Neural networks are arranged in a branching pipeline, with rules coming from domain knowledge used in between to reduce the complexity of the problem and make it more tractable. Later steps utilize knowledge extracted from previous steps to facilitate further processing, significantly improving accuracy compared to the naive, direct approach. This divide-and-conquer approach is possible thanks to our annotation of data into a hierarchical scheme, which lets us partition the problem into a series of sub-tasks, each easier than the full task.

One difficulty of our research lies in the sparseness and complexity of the data. There are many different symbols to learn to predict, some with different graphical variations, and some of which can be rotated or mirrored to convey different sign language meanings. These symbols are arranged in a single SignWriting transcription in a meaningful way, so it is necessary to find their relative positions at the same time as the symbol meaning is found. Combined with the small size of the available dataset, this complexity makes the search space of our problem too sparse to solve with direct application of existing neural networks. By reducing the features to extract at each step, each search space is made more dense, and entropy of the problem reduced. Not only are each of the sub-tasks then easier to solve, but the information extracted can be further processed, informing the next step and improving its results.

To recapitulate, in this article we present a novel problem for the artificial intelligence community: automatic recognition of SignWriting instances. We propose a modelization of the problem, by giving a computationally tractable description of the underlying data, and we offer a solution using some common machine learning algorithms. Due to the originality and inherent difficulty of the problem, our solution combines the machine learning algorithms with rules coming from domain knowledge. We validate our approach by evaluating its performance in solving the task, and also compare it to the performance of a direct, single algorithm, one-shot approach. These first solutions to the novel problem

of SignWriting recognition validate our computational modelization of the data domain, and present an opportunity for future research to improve on the task.

The rest of the paper is structured as follows: in section II, we present an overview of Sign Language and SignWriting, to show the complexity of the problem and introduce the key points necessary to understand our system. Section III is devoted to other solutions that try to solve similar problems, and briefly presents useful deep learning approaches. Our contribution is divided into two sections: section IV discusses the underlying data engineering, and section V presents our solution as well as the baseline approach. These are evaluated and compared in section VI, and finally, section VII gives our conclusions and outlines future lines of research and development that can be followed.

II. SIGN LANGUAGE AND SIGNWRITING

Sign languages, as the natural languages of the Deaf and Hard-of-Hearing communities, utilize not sound, but vision. Sound, as used in oral speech, is a wave of air pressure modulated in time to create different qualities to which we then assign meaning. These sounds are easily quantized into individual units, the phonemes, so the fundamental model of speech in linguistics and language technology is as a sequence of individual phonemes.

This simple model is then easily transferable to writing, where phonemes are represented by characters. To each individual sound, we assign a character (sometimes more), and then write these characters in order. Instead of phonemes, some languages assign characters to syllables, morphemes, or whole words, but in their fundamental, written texts are strings of characters each representing discrete sounds or sound sequences.²

In sign languages, it is not clear what the fundamental unit is, and this is quite an actual issue of linguistic research. What constitutes phonemes, syllables or even words in sign language? There is a tension in linguistics between trying to apply existing categories, developed from oral language research, and recognizing the uniqueness of sign language characteristics, but there are many examples of in-depth descriptions of different sign languages' grammar and phonology in the literature [14]–[17];

Some systems developed to write sign languages, such as that devised by [8], or the Hamburg Notation System [6], are in their origin linguistic notation systems, but SignWriting [9] is a naturalistic system based on iconicity rather than linguistic categories. As such, the fundamental “unit” in SignWriting is the hand, since it is the most prominent articulator in the phonology. We will call the symbols used in SignWriting “graphemes” from now on, to highlight their nature as units of a writing system, but different from the linear “characters” of oral writing systems. We will introduce

²“Ideographic” writing systems, such as the Han characters in use in Chinese and Japanese, seem different on the surface, but in reality each character is assigned a set of possible sounds, so the sequential, one dimensional model still applies.




FIGURE 2. Some hand graphemes, comprising different hand shapes and orientations. Based on [19].

the different SignWriting graphemes in the following, to give a sense of the problem we are trying to solve, and use them to introduce the relevant parts of sign language phonology which give rise to the complexity of its writing systems.

Since hands are the most prominent and main articulator of sign languages, so hand graphemes are the most visible graphemes in SignWriting. They are iconic depictions of hands, using polygons to represent their configuration, namely the shape the palm and fingers make. The fingers can flex against the palm, extend, join laterally, curl, or even cross. The possibilities are somewhat different in each sign language, but see [18] for an analytical account and coding system for those in American Sign Language. Hand graphemes can be used to represent right or left hands, which are mirror images of each other, so the shape is horizontally flipped (reflected) to represent them when the hand grapheme is not symmetric.

As objects in three dimensional space, hands have a location and rotation which, when relevant, need to be specified.

First, orientation captures the rotation of the hand in three dimensional space. The same sign, with hand rotated in different angles, can mean different things, so orientation is an essential parameter to notate. However, it is an intrinsically three-dimensional feature (see our somewhat mathematical account in [19]), so a number of graphical techniques are required to depict it in the flat surface that SignWriting uses, enumerated in the following and depicted in Fig. 2.

1. To project the 3D hand into a 2D plane, a plane of observation has to be determined. This can be the vertical plane (parallel to a wall in front) or the horizontal plane (parallel to the floor).
2. If the chosen projective plane is the floor plane, the fingers are drawn separate to the hand, like in the last example of Fig. 2.
3. The hand grapheme can then be rotated, to iconically represent the rotation as observed from the chosen




FIGURE 3. Examples of non-rotating graphemes. On the left, some body parts (on the head) are shown. Clockwise from top left: generic face, eyes, nose and ears. On the right, diacritics marking different types of contact are shown. Clockwise from top left: touch, rub, brush and strike.




FIGURE 4. A selection of movement markers in SignWriting. Each of these is a different character in SignWriting fonts, but we can think of it as composed of smaller units: straight or curved segments, single or double stems, and black, white or empty arrow heads.

point of view.

4. To indicate wrist rotation, the palm of the grapheme is filled in with different patterns according to how it would be seen from the chosen plane. If the palm is facing toward the signer, the grapheme is white, while if the back is seen, it is filled with black. If the hand is partially rotated, it is filled half white half black, with the white side indicating where the palm is.
5. To improve iconicity, when the different fills are used, the positioning of the fingers in the grapheme can change, to more iconically capture the observed shape of the hand. This can mean that there is sometimes ambiguity between left and right hands, when their wrists are rotated such as the filling ends being the same.

Next, the hand has to be located in space, since a sign performed at the height of the head is not the same as the same movements performed at the chest, for example. In SignWriting, a number of graphemes exist to denote different parts of the body. These graphemes can then be placed in the page, and hand graphemes located relative to them to iconically represent their three-dimensional location. If the sign is realized in the “neutral” space (in front of the body, but not relative to any particular body part) no body graphemes need to be used, or if the head is required³, the hands are placed sufficiently distant to it to not give rise to confusion. Of course, the exact placement of hand and body graphemes is a subjective and stylistic decision, which may not often be a problem for humans but can be an obstacle for computational processing.

Related to location, an important feature in sign language is contact. Hands can touch, rub, brush and strike each other or different parts of the body. This is represented using small graphemes, which we call diacritics, placed in suggestive positions near the point of contact. These are the main diacritics, but we also group under this umbrella other small, independent graphemes which capture things like movement dynamics, pauses, or internal hand movements. Internal hand movements capture an evolution of the hand shape in the sign, such as fingers being bent, extended, or wiggling. Some

³The head may be required to show the eyes or mouth, which can alter a sign’s meaning, for example with a frown or a smile.

examples of heads and diacritics can be seen in Fig. 3.

Lastly, hands are often not static, but rather move in space, and these movements can have meaning in themselves. In SignWriting, movements are represented with arrows, which can have a start and an end, and depict a straight, curved, or more complicated trajectory between those points. Movements can also be circular, naturally represented with a closed circular trajectory. To represent three dimensional movements in the flat page, the same observation planes used for hand orientation are required. When the movement occurs in the plane parallel to the floor, single-stemmed arrows are used, while double-stemmed arrows represent movements in the vertical plane. Additionally, the arrow heads can be black, if the right hand is moving, filled with white, if it is the left hand moving, or left open, when both hands move together as a unit. Movements can get very complicated, and some examples are presented in Fig. 4.

These are not all of the possibilities of SignWriting, and even within the described graphemes there is complexity that is out of scope for this article. This introduction should be enough to understand the complexity of the problem in hand, and the reasons for our proposed solution. Nonetheless, a complete textbook by the inventor can be found in [20], or online at <https://www.signwriting.org/>.

III. COMPUTATIONAL APPROACHES

From a language technology point of view, one might think of our problem as a task of Optical Character Recognition (OCR). OCR is the task of recognizing the written form of language found in images, often scanned, and either coming from a printer, other type of digital production or handwritten. Indeed, that is exactly our problem, recognizing text from rasterized images where the language information is lost but only pixels remain. However, existing OCR solutions do not work for our problem. [21] gives a fascinating account of the history and architecture of Tesseract, an open source OCR engine, but the reader will note that most of the issues, tricks and clever solutions are based on an understanding coming from oral language.

Like Tesseract, existing OCR systems are developed for the writing systems of oral languages, and thus often rely on two assumptions that don’t hold for SignWriting, both of them related to the spatial uniqueness of SignWriting–

which is in turn an effect of the spatial use of the viso-gestual modality of sign languages.

First, graphemes are assumed to form one-dimensional sequences. Text flows in a main axis, sometimes wrapping along a second axis (e.g. in western languages, characters flow from left to right and then lines can wrap from top to bottom). Even with the more advanced OCR systems that can detect text in different, random positions in the image, when characters form a word they are expected to form a line. In SignWriting, graphemes are distributed spatially in a significant but non-linear way.

Related, graphemes in oral language writing systems are expected to be mostly upright. Even some typographical variations like cursive don't rotate the character more than a few degrees, and in any case this rotation is not significant but rather stylistic variation or noise. In SignWriting, hand graphemes and movement markers rotate and reflect to represent different spatial meanings. While this could be solved by treating each of the possible transformations as a different grapheme, this multiplies the number of graphemes to recognize in more than an order of magnitude. This makes it unpractical both for the manual processing and tagging required to get the expert system running, and rarefies the data, making it sparser and therefore the use of deep learning less robust.

A similar problem to ours is faced by [22], who need to recognize a different kind of language: handwritten logical circuit diagrams. In those diagrams, the 2D position of circuit elements is meaningful, and so, as for us, OCR algorithms are not applicable. However, they have a much smaller vocabulary than ours. SignWriting graphemes number in the hundreds, while logic circuit elements in their dataset seem to be less than ten. Additionally, they only deal with their elements in a standard orientation, the circuit flowing from left to right, so they do not have the problem of rotations and reflections of the graphemes.

Indeed, with the recent advances in artificial intelligence and deep learning, more such graphical writing representations may become apparent⁴. Since ready-made OCR systems cannot be used, nor OCR technologies be adapted to the problem, it is necessary to directly use the underlying technology: computer vision.

A. COMPUTER VISION

Computationally understanding images is a hard problem due to a number of factors. Of course, what humans see and interpret in an image is a rich and subjective composition of different meanings, but we often just need a simpler understanding, such as labeling the object depicted (classification), finding different objects in a scene (object detection, scene understanding) or separating image regions according to the real life object to which they pertain (image segmentation).

The main issue is that the way images are usually stored and manipulated is completely unrelated to the way humans

understand them. Instead, they are optimized for display on square arrays of color elements (monitors and screens) and are therefore stored as arrays of pixel values. If an image is stored in row-order, neighbour pixels in the vertical direction will be far apart in memory. If an image stores the different color information (channels) separately, even the values needed to reconstruct a pixel will be far apart in the computational representation.

This pixel array-based representation thus presents a problem for algorithms that try to extract meaning from images, since it is unfeasible to write deterministic and exhaustive rules that relate when pixels form lines, shapes, or more complex objects.

To deal with this, techniques which compute aggregated features from the pixel information have been used to great success. Some of these act globally, computing mathematical properties of full images, while kernel methods use convolutions to compute local properties of images, by integrating with a kernel function suited to the particular task which takes into account the values of nearby pixels. Many different mathematical algorithms and statistical techniques have been developed, and a comprehensive overview of the state of the art before the uprise of deep learning can be read in [24].

But in recent years, there has been an exceptional expansion of machine learning techniques, especially around the use of deep learning, which has notably improved the state of the art both in accuracy and range of problems which can be tackled. While neural networks haven't necessarily replaced traditional methods (see [25] for a discussion on this), they have become very popular, not only for their success rate but also because of their relative ease of use.

A rough description of machine learning techniques is to use training data, a large number of input examples where we know the desired output result, to decide the parameters in a regressive (predictive) algorithm by minimizing the error made. For example, in the most common neural networks, an iterative process of prediction and error computation (forward and backward propagation) is performed and the algorithm parameters are iteratively improved.

The strong suit of neural networks is their ability to extract and remember patterns in source data, without the need for the researcher to accurately formalize or describe these patterns. Deep neural architectures can build these patterns from other patterns, in a cascade of increasing complexity, which makes them particularly suited to computer vision. One can imagine pixel arrays turning into lines and shapes, lines and shapes into body parts, and these body parts being then combined to discover that an image is that of a dog.

B. NEURAL ARCHITECTURES

As we said at the start of the section, there are different tasks in computer vision, and neural networks are not only applied to computer vision, but rather a wide variety of problems. All these networks are not the same, but each problem requires a specific architecture (combination of layers, activation func-

⁴We discuss some more of this in [23].

tions, and other parameters) suited to learning the particular patterns that solve it.

For image classification, that is, finding a suitable label to describe the content of an image, a usual neural architecture used is that of convolutional networks, for example AlexNet [26]. This architecture is composed of convolutional layers, which combine local features of an image into aggregated characteristics. These patterns are built upon, layer after layer, and in the end, sufficiently sophisticated graphical characteristics are found, so a probability of the original image belonging to a particular class can be given. As in all neural networks, the patterns and convolutions found by the network are not pre-determined, but rather optimized in a training step using already annotated images.

“You Only Look Once” networks (YOLO) are a type of neural network used for object detection and classification. Given an input image, the different objects contained are found, and a label assigned to each of them. This architecture can be trained for different corpora, and is described in [27]. Roughly, YOLO networks learn, for the different regions in the image, the probabilities that a particular object or its boundary can be found there, and then reconstruct the objects’ positions and sizes from this information.

IV. DATA DESIGN

In the previous section we have seen machine learning approaches to the task of computer vision. But the quality of any machine learning approach rests on a sometimes overlooked cornerstone: data. Training data must be sourced, prepared, preprocessed, and often annotated. This is a costly process in terms of time and effort, but on its correctness lies the maximum real performance of the trained algorithms in the real world. As [28] say, an error in any step of the data pipeline propagates to further steps, compromising the accuracy and reliability of the algorithm. When approaching a novel problem, often new data must be sourced. Furthermore, the problem must be modeled, and the data annotation schema designed. Neural networks learn patterns very accurately, but we have to decide what patterns to learn, what features to discriminate and which to abstract.

This process of data engineering is a fundamental step of the expert resolution of a problem, a step where much of the domain knowledge to be used is embedded in the final system, and thus it is as much part of the solution as the system’s code and implementation.

Unfortunately, there is not a lot of SignWriting data available to us, and what can be found is not in an easily processable format. Sometimes there is an association of transcription to meaning in oral language (which doesn’t really help, since the oral word is almost never related to the sign parameters) and most of the time there is no indication whatsoever of the symbols contained in the SignWriting transcription.

This has meant that we have had to prepare our own corpus of SignWriting that can be used for machine learning or linguistic research, using data collected in a collaboration




FIGURE 5. Annotation for a hand grapheme, showing the different labels and values that need to be assigned. To its right, an illustration of the depicted hand shape is shown.

with linguists as part of a project to develop tools for the ease of SignWriting visualization: <https://www.ucm.es/visse>.

In the corpus, there are two types of data: logograms and graphemes. Logograms are full transcriptions, images of SignWriting which correspond to a sign. Graphemes is the name we have chosen for individual graphic components, units of linguistic information which by themselves convey some of the meaning necessary to reconstruct the sign. The full information is obtained by combining the meanings of the different graphemes while taking into account their relative position and rotation within the logogram.

The information represented by each grapheme, however, is itself complex, and tagging each of them with a single label is not enough. We want to be able to discern the different features described in section II, and so we have developed an annotation schema for our corpus, where we assign a variable number of labels to each grapheme, depending on how much information is needed.

A first label, “CLASS”, divides graphemes into 6 classes of symbols, with related graphical properties and similar semantic information: “HEAD” for head symbols, “DIAC” for small diacritics, often used for contact information; “HAND” for graphemes which represent hands, and “ARRO”, “STEM” and “ARC” for arrows.

Graphemes are then annotated with an additional tag, “SHAPE”, which identifies the actual SignWriting symbol. For HEAD and DIAC graphemes, this is enough information, but other graphemes require more annotation. HAND graphemes, such as the one tagged in Fig. 5, have an additional label, “VAR”, which represents the wrist rotation (whether the hand is white, black, or half and half, see section II about orientation). With this, the “lexical” information of a hand grapheme is fully specified, but we still need to record the rest of its orientation, as graphically represented by the symbol rotation and mirroring. Rotation (ROT) can take one of 8 different values, described with the cardinal points (N for north, NE for northeast, etc.). If the symbol is mirrored (horizontally flipped) a value of ‘y’ is added for the reflection (REF).

The remaining classes of graphemes, ARRO, STEM and ARC, are those needed to represent movements. As we saw in Fig. 4, movement markers in SignWriting can be very complex, as well as mixed and superposed to create compositional meaning. Nonetheless, they can be seen as composed

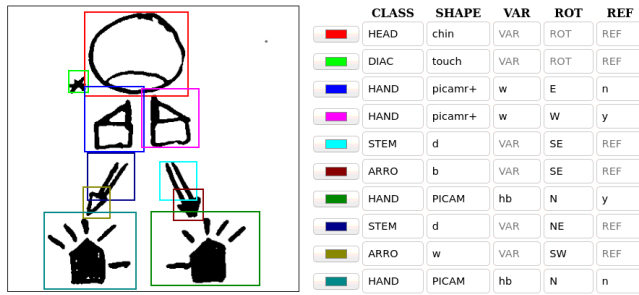


FIGURE 6. Annotation of the logogram for the sign “History”.

of different segments: the straight or curved trajectories, or the end of movement arrow heads. Even if, from the point of view of SignWriting, movement markers are single holistic “symbols”, each of the segments has its own visual identity and meaning, so in our corpus we have chosen to identify them as separate graphemes.

Arrow heads (class ARRO) can be black, white, or not filled, to represent which hand is moving. The body of the arrow can then be straight (STEM) or curved (ARC). STEMS and ARCS can be either single or double, to represent whether the movement is parallel to the floor or to the wall, and finally ARCS can be a full circle, half of one, or a quarter. This is all encoded into the SHAPE, and it is our hypothesis that with this graphemes most if not all movement markers can be recorded. Finally, graphemes of ARRO, STEM and ARC class can rotate, so need an additional ROT tag.

We have collected and annotated 982 handwritten logograms, within which 6071 graphemes can be found. The logogram annotation consists of the bounding boxes (location and size) of the graphemes, and then tags for each grapheme with the appropriate combination of features. An example annotation can be seen in Fig. 6. The dataset we have created is publicly available at <https://zenodo.org/record/6337885>, including an annotation guide with detailed information about the different grapheme classes and their tags.

V. PROPOSED SOLUTION

In the previous section, we have described the data that we have available, and an annotation schema with which to organize and store the information contained within each image in our dataset. This schema is a modelization of the problem we want to solve: extracting a computational representation of the semantics represented by SignWriting images. As we saw in section III, we can train machine learning algorithms to learn the underlying patterns in annotated data, so that they are able to reproduce them in other, previously unseen, data samples. Therefore, we can use these algorithms, in particular neural networks, to learn and reproduce the modelization of our problem implicit in our corpus annotations, and solve our problem this way.

As a first approach and baseline for comparison, we use a state-of-the-art computer vision algorithm: YOLO. This

is a deep learning detector and classifier neural network architecture, which can solve the full task with one single network, but has many limitations. Our proposed solution is a second approach which addresses some limitations of the first, combining additional networks into a branching pipeline which intelligently decides the paths to follow. This allows us to exploit the expert knowledge encoded into the corpus annotations, and solve a data problem in which a Big Data approach is not enough.

A. ONE-SHOT APPROACH

YOLO networks, as introduced in section III-B, seem like a perfect fit for our problem. They solve detection and classification in one step, so given a logogram, they should be able to locate the different graphemes depicted and assign a label to each of them. A problem might be that our graphemes are tagged with many different features, but we can create a different label for each feature combination (concatenating tag values) and give this resulting label to the network to learn. From this single label, the different features can then be extracted, and so the full task of SignWriting resolved. We use YOLO version 3, from Redmon and Farhadi [29], as implemented in <https://github.com/AlexeyAB/darknet>.

We will examine the results later, in section VI, but this one-shot solution will prove to be not enough for our problem. Essentially, there is too much complexity in our data for the single neural network to solve it successfully.

B. EXPERT KNOWLEDGE-BASED PIPELINE APPROACH

To make the detection problem more approachable, we have to reduce the information that is passed to the neural network. Instead of telling it that each possible different grapheme is an object to locate, we can omit some of this information, and just have the network learn to find a few rough groups of graphemes (the CLASS tag from section IV). This way, the detector network learns to find objects abstracting away some of their details, which makes it more able to generalize, and thus increasing its ability to find graphemes which are not seen exactly so in the training data.

Then, we can use a different procedure to complete the missing features. Since detection has already been performed, we can extract the region of the image where the detected grapheme is found, and utilize only this sub-image. This lets us ignore all information unrelated to the particular grapheme in question, and what we have left is a new task of, given an image depicting just a grapheme, find out the features it represents. Moreover, since we already have the rough grouping given by the detection step (the CLASS tag), we can use different processes for each of the grapheme classes.

For most of them, it is enough to use a single neural network to learn all of the remaining features. We use AlexNet, as introduced in section III-B, and treat this problem as one of image classification. For each grapheme, a label is computed by concatenating its features, just like was done for the one-shot YOLO. The difference is that now we train different

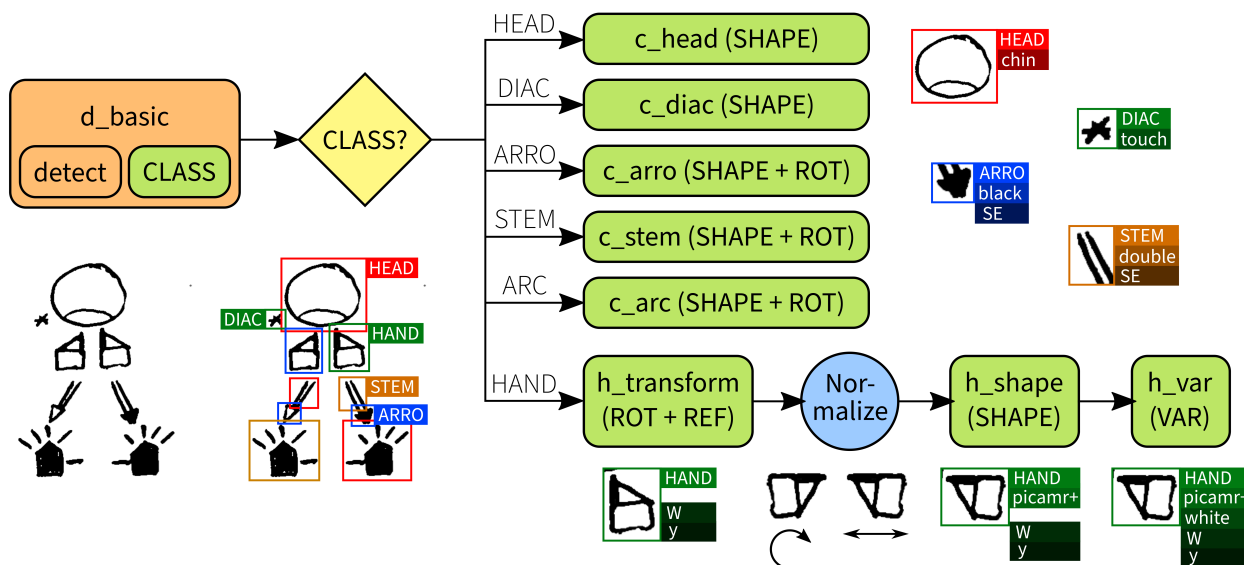


FIGURE 7. Full pipeline architecture, including the different networks, decisions and processes. An example logogram and some of the detected graphemes are shown as they progress through the pipeline.

networks for each of the grapheme classes, reducing the number of labels they have to learn, and letting them concentrate in the specifics of each different grapheme instead of the commonalities to the whole group.

The case of hand graphemes is a little different. As we saw in section IV, hand graphemes have each four distinct features in addition to the CLASS, that is the SHAPE (finger configuration), the VAR (palm orientation) as well as the graphical transformation (ROTation and REFlection). There are more than 50 hand shapes in our corpus, which multiplied by a possible 6 different variations, 8 rotations and 2 reflections (normal and mirrored) give a total of more than 4000 possible combinations. This is too difficult for the network to learn, more so given the scarcity of our data, so we further subdivide the problem.

First, a classifier (again AlexNet) is used to detect the transformation of the grapheme (rotation and reflection). Once the transformation is detected, the grapheme is normalized to a standard form (North rotation, no mirroring). A second AlexNet classifies this normalized form, finding its SHAPE, and a third and last network finds the VAR. This way, we encode our knowledge of SignWriting into deterministic rules and steps of processing. The neural networks are still in charge of detecting the visual patterns that constitute the graphemes, a task better suited to deep learning, but reconstructing the grapheme tags from the different visual patterns is done in the glue between the networks. This means each network only has a smaller task to learn, which it can tackle even with the limited data available. The normalization of rotation and reflection is again a deterministic rule, which transforms the data in ways that we know make sense, freeing the networks from having to do the work of generalizing such transformations.

The full pipeline for SignWriting can be seen in Fig. 7, with an example logogram for reference.

VI. EVALUATION

There are many different metrics which can be used in machine learning and computer vision, including precision, recall, mean average precision, etc. These metrics are each useful for different things, and for a complete understanding of an algorithm it is often necessary to measure all of them, and obtain a rounded view of the problem. Indeed, we have used them while building our system. However, we are not so much interested in the performance of the networks themselves, since there is extensive research on this, but rather in the performance of our system for our full problem, that is, transformation of SignWriting images into useful computational representations.

To measure this, we use accuracy, which is a balanced scoring metric and has the added benefit that it can be directly compared between our solution and the baseline direct approach. Accuracy is a measure of how good predictions by an inference system are, by dividing correct predictions by the total number of instances. We adapt the concept (proportion of correct over total) to three different measures which are useful to us. The key aspect to evaluate is the predicted graphemes: how many of them can be found, and how accurate their predicted features are. The third measurement combines the first two into an overall performance of the full task.

First, we need to measure detection accuracy, that is, whether predicted graphemes are actually there, and whether the graphemes which are there are found at all (equation 1). A grapheme is considered correctly predicted if the area of the bounding box sufficiently overlaps the true bounding box.

If there is no true grapheme where a predicted grapheme is found, it is not counted as a correct detection, and if no grapheme is predicted where the true annotation has one, it is counted as a missed detection.

$$\text{Detection Accuracy} = \frac{\text{correct detections}}{\text{total detections} + \text{missed det.}} \quad (1)$$

Then we need to measure whether the labels assigned to each grapheme are correct, that is, classification accuracy. This is only measured for graphemes which are correctly detected, and it is the proportion of the correctly classified graphemes over the total number of correct detections (equation 2).

$$\text{Classification Accuracy} = \frac{\text{correctly classified}}{\text{correct detections}} \quad (2)$$

Finally, a combined measure is computed, overall accuracy (equation 3), which scores each solution for the global task of recognizing SignWriting. It counts the graphemes correctly detected and classified, dividing it among the total number of predictions and missed predictions.

$$\text{Overall Accuracy} = \frac{\text{correctly classified}}{\text{total detections} + \text{missed det.}} \quad (3)$$

Accuracy measures are proportions, so the three given scores range from 0 (worse) to 1 (perfect) values.

To evaluate the performance of the algorithms closer to how they would be used in the real world, and as is standard practice, we split our logogram data into two sets: a training set, used to automatically learn the patterns, and a testing set which is not used in training. This set is never seen by the algorithm, and thus simulates real world, previously unseen, data. The different accuracy measurements are evaluated on this test set. There are 791 logograms in the training set, and 191 in the test set, resulting in a total of 4840 graphemes in the training set and 1231 in the test one.

The results of computing these metrics for our solution and the baseline single YOLO network are shown in Table 1. The overall accuracy of the direct approach is of 0.58, which may seem low but is impressive for the complexity of the problem. Our pipeline manages to increase performance by 17%, to an overall accuracy of 0.68. This improvement alone shows the validity of our approach, and a human analysis of the pipeline results gives an even more optimistic view. Often, the pipeline makes mistakes that are less severe than a complete failure. For example, similar hand shapes are often confused, or diacritics mixed up. While these are indeed wrong predictions, and are counted as such in the accuracy computations, the partial truth that they are able to predict can still be useful for downstream applications to process, and this is the great advantage of the pipeline.

Our full experimental setup can be reproduced using our published dataset (<https://zenodo.org/record/6337885>) [30], which includes the Data Version Control (DVC [31]) configuration files that define the experiments, and scripts for performing every step. These scripts use our software Quevedo [23], a tool for annotating and processing graphical language

TABLE 1. Performance of our solution and a baseline one-shot algorithm for the task of SignWriting recognition: grapheme detection and classification within logograms.

Solution	Detection	Classification	Overall
Single YOLO	0.67	0.88	0.58
Pipeline	0.86	0.78	0.68

datasets. To examine the repository, for example one could use the following commands:

```
$ wget https://zenodo.org/record/6337885/
  files/visse-corpus-2.0.0.tgz?download=1
$ tar xzf visse-corpus-2.0.0.tgz
$ cd visse-corpus
$ pip install quevedo[web]
$ quevedo web
```

To reproduce the machine learning evaluation, one also needs to have DVC and Darknet installed. Then, it is as simple as issuing the DVC “reproduce” command:

```
$ dvc repro
```

This includes all information needed to reproduce our experiments, since all the steps and data of our experimental setup are contained in the different configuration files and software. However, since Quevedo is generic, not only can our SignWriting experimentation be reproduced, but also its ideas applied to other datasets and domains.

A. ERROR ANALYSIS

In a more in-depth analysis, the first immediate observation is that the single YOLO detection performance is too low to be useful. While its classification score is good, this is an effect of only classifying a handful of graphemes, the ones that have been detected. Detection, however, misses too many graphemes, essentially ignoring those which are not common enough. While focusing on the most common data is not a bad strategy in many situations, in this case detection performance is too low to justify it. Furthermore, for our purposes, incomplete predictions can be useful, since a lot of the meaning in the transcription can be later reconstructed, which can not be done for a missed detection.

As was advanced before, the probable reason for this lower detection accuracy is that, by giving all the features to the YOLO algorithm, it can not see the common properties of the different grapheme classes. We tell it that a touch diacritic is a different thing than a rub diacritic, so it needs to differentiate them and can not exploit their graphical similarities. This impedes proper generalization, and thus the network only learns to detect and classify graphemes which it has seen enough, ignoring the rest.

In our pipeline solution, only the CLASS feature is given to the detector network to predict. The different grapheme classes have been chosen due to their graphical properties, so the network can exploit this to learn to discriminate them while at the same time being able to generalize to instances

TABLE 2. Detection, classification and overall accuracy of our pipeline solution and the baseline one-shot algorithm, computed for each grapheme CLASS.

CLASS	One-shot			Pipeline		
	Det.	Class.	Overall	Det.	Class.	Overall
ARRO	0.79	0.93	0.73	0.90	0.85	0.76
HEAD	0.80	0.81	0.65	1.00	0.65	0.65
DIAC	0.82	0.89	0.73	0.90	0.83	0.75
STEM	0.69	0.91	0.62	0.83	0.89	0.73
HAND	0.48	0.74	0.35	0.88	0.61	0.54
ARC	0.51	0.56	0.28	0.44	0.55	0.24

not seen before. In fact, the YOLO network is better at this rough classification than a grapheme classifier network trained specifically for this task. It is likely that in this case, having the full logogram context can help, rather than hinder, prediction. Diacritics are smaller than hands, which are smaller than head graphemes. Arrow components (heads, stems and arcs) are usually found together. This context is lost when isolating the grapheme, but present in the full logogram, so the detector can use it to give us a first rough classification which we then apply to split further processing into branches.

Regarding topographic accuracy of detection, that is, how close the predicted bounding boxes are to the annotated ones, it is generally good across different configurations that we have tried. Detecting and separating black-on-white objects is generally easy for the network, with two very relevant exceptions. The first problem are diagonal elongated objects, that is, arrow stems. These graphemes are sometimes very long, and when rotated, they may actually occupy very little of their bounding box—the bounding box is square, but the stems only fill the diagonal. This can be a problem when other objects are present in the bounding box, even if not overlapping the actual arrow.

The second problem is, precisely, overlaps. While YOLO networks seem to do a good job with partially obstructed objects, sometimes graphemes are placed in a “cross” configuration, where they overlap diagonally, and the ends of the lower grapheme spread to both sides of the overlapping one. To further complicate this issue, graphemes so placed often have the same CLASS. Hands can be placed on top of each other, movements can have cross-like trajectories, etc. It is likely that YOLO networks have trouble with these combinations due to the way the edge and interior probabilities are merged by the network to find the predicted bounding boxes.

Fortunately, while not uncommon, these two detection problems are the only issues in this step, and do not hinder further classification of graphemes.

By splitting the accuracy measurement for each grapheme class as in Table 2, more detail can be seen. It is clear that two grapheme classes are especially problematic: HAND and ARC. Their classification is a tough problem, which can be

seen in the low detection of the single shot YOLO network (remember that it performs detection and classification at the same step, so difficult to remember graphemes will not be detected at all) and in the classification accuracy of the pipeline.

ARC graphemes, which represent curved movement trajectories, are numerous and very varied, while at the same time being the class with lowest number of instances found in the corpus. It is also the case that hand-drawn arcs tend to present the highest graphical variability, since many different angles and sizes can be used by the transcriber to represent the same meaning. Both issues lead us to think that increasing the number of ARC instances in the training data is the most important step to be taken, but more detailed processing like the one done for HANDs may also help.

HAND graphemes are also very difficult, probably more so than ARCs, due to the multitude of features to be predicted and how they interact. However, hands are probably the most prominent articulator of sign languages, and as such appear in good numbers in the corpus. This can be seen in the huge leap in detection accuracy, from 0.48 by the single YOLO network to 0.88 by the pipeline, and the overall accuracy improvement from 0.35 to 0.54.

In both cases, classification accuracy affects overall performance, but detection accuracy is much better with the pipeline than in the YOLO single shot solution. As said before, being able to detect that a grapheme is present is fundamental for correct computational representation of SignWriting. It is of utmost importance to detect hands, which the single YOLO often fails to do, and in the case of ARCs knowing that a sign has some curved or circular movement, even if the concrete details are not known, is already very useful information.

B. LIMITATIONS OF THE SYSTEM

There are a number of limitations to our work, the most important of which is related to the data used for training the deep learning algorithms. The corpus we have collected is very representative of the different graphemes available for composing SignWriting, but there are other dimensions along which there is not so much variation. For example, most of our transcriptions come from a single informant, so there can be stylistic choices made by that informant and peculiar to their use of SignWriting. Additionally, our logograms represent signs from Spanish Sign Language, so graphemes that codify features not common in Spanish Sign Language may not be properly represented. To overcome this limitation, more data need to be available, and we will work in the future to acquire them.

Thinking about the graphical context of the logograms, all of our samples are in black font on white background, and have been preprocessed to maximize contrast. This can be solved with more data, from more contexts (e.g. ruled or squared paper, using colored ink, from a camera photo with bad lightning) or by adding a preprocessing step to our system which corrects for this kind of variations.

As to the machine learning algorithms employed, it is likely that there are better algorithms available. The state of the art in deep learning is rapidly improving, and it seems every other year there is a major breakthrough. This algorithms could be swapped in, replacing the neural networks we use and improving accuracy. However, since the difficulty of our problem lies in the complexity of the data and the small amount available, our system can continue to work as is, augmenting the machine learning techniques with the expert knowledge and therefore improving overall performance.

C. PRACTICAL APPLICATION

The system described in this article, including the featureful description of SignWriting instances and the automatic pipeline capable of extracting it, underlies the “Visualizing SignWriting” application (<https://github.com/agarsev/visse-app>). This application was one of the results of project “Visualizando la SignoEscritura” (VisSE, “Visualizing SignWriting” in Spanish) which sought to create tools for the better use of SignWriting in the digital world. The user-facing result is a web application able to recognize instances of SignWriting, be them scanned from the user device or created from an image processing software, and explain them to the user via textual descriptions of the graphemes and 3D models of the hands. This application validates our approach by showing, on the one hand, the use of image recognition to process SignWriting, and on the other hand, the usefulness of our annotation schema which can be leveraged to generate explanations for the components of a newly seen logogram.

VII. CONCLUSIONS AND FUTURE WORK

As the importance of sign languages is recognized throughout the world, the inclusion of the signing community in the digital economy is fundamental. To this end, computational representations of sign languages are necessary, both for end users and researchers alike. In the case of SignWriting, much of the available data exist in image formats, understandable by humans but not machines. Converting SignWriting images into a computational representation is a necessary first step to automatically process them, but requires state-of-the-art applications of artificial intelligence, since images are not easy to process using hand-crafted rules or ad-hoc procedures.

We have presented a careful analysis of the data underlying the problem, establishing a categorization of the different meanings of SignWriting symbols into hierarchical features. This formalization is itself one of our contributions, a computationally valid representation which captures the intended meaning of SignWriting transcriptions into a numerical representation of the positions of graphemes within a logogram and additional key-value pairs of features for each of them.

To automatically reproduce this representation for new instances of SignWriting, the best approach is to use deep learning, able to capture complex relationships in the data and generalize from the patterns present in a training corpus to the general case. However, large amounts of data are necessary to make deep learning approaches work reliably.

For our problem, there does not exist a reference corpus of data, or a similar problem from which to transfer learned neural network weights.

We have collected the necessary samples, and created a corpus [30] with which to train the algorithms. Still, the amount of data available is small, and costly to annotate. However, we have shown that the annotation pays off if done carefully. Our use of many features, decided both from a semantic point of view and the necessities of the visual processing required for the problem, has allowed us to build an expert solution able to automatically recognize SignWriting. Compared to the simple, direct approach using a single YOLO network, our proposed improved system uses many deep learning networks, combined in an intelligent pipeline which can extract additional information and make decisions based on previous steps of processing, achieving a 17% improvement in recognition accuracy and additionally being able to extract partial information even when recognition fails.

On the whole, domain knowledge about the problem has let us create a system which utilizes deep learning approaches even in a situation where no existing data can be found, by collecting the corpus ourselves, defining a formal schema for its annotation, and exploiting it to get the best performance from the neural networks employed. Our ideas and approach may be useful not only to process SignWriting instances, but may be applicable to other problems where the data available are less numerous than the expert knowledge that can be collected.

A. FUTURE WORK

There are three straight-forward directions in which this research can be improved. On one hand, collecting more and more varied data will likely improve performance of the system, and solve some of the limitations outlined above. A second direction to go is downstream, putting the recognized SignWriting representation to use in more consumer applications. The needs of these applications will tell us what the strong points of our approach are, and where its need to improve to support their use case. Finally, the components themselves used in the system may be improved. We have used readily available neural network architectures, as can be found in the literature, and with implementations that we can directly use. Fine-tuning the network parameters, or swapping some of them for networks better suited to each particular sub-task, will surely improve overall performance.

There is also room for alternative approaches to be tried. An ensemble of neural networks, where their results are weighed and combined, can help improve detection and classification of rarer graphemes, or correct frequent errors for certain common or uncommon situations. A custom neural architecture that embeds all the steps in our pipeline may be possible, which would facilitate feedback between steps, or some other technique for improving the results of earlier steps in the pipeline by taking into account the confidence of later steps.

REFERENCES

[1] G. Barberà Altimira, *The meaning of space in sign language: reference, specificity and structure in Catalan sign language discourse* (Sign languages and deaf communities 4). Boston: De Gruyter Mouton : Ishara Press, 2015, ISBN: 978-1-61451-866-2.

[2] D. Brentari, *Sign Language Phonology*, 2019.

[3] S. J. Parkhurst and D. D. Parkhurst, *Variación de las Lenguas de Signos usadas en España* (Revista Española de Lingüística de las Lenguas de Signos), trans. by A. P. Sobrino. 2001.

[4] R. J. Senghas, A. Senghas, and J. E. Pyers, "The emergence of Nicaraguan Sign Language: Questions of development, acquisition, and evolution," *Biology and knowledge revisited: From neurogenesis to psychogenesis*, pp. 287–306, 2005, Publisher: Lawrence Erlbaum Associates Mahwah, NJ.

[5] W. Sandler, M. Aronoff, C. Padden, and I. Meir, "Language Emergence: Al-Sayyid Bedouin Sign Language," in *Cambridge handbook of linguistic anthropology*, ser. Cambridge Handbooks of Linguistics, Oct. 2014, pp. 250–284, ISBN: 978-1-139-34287-2.

[6] T. Hanke, "HamNoSys – Representing Sign Language Data in Language Resources and Language Processing Contexts," in *Proceedings of the Workshop on Representation and Processing of Sign Language, Workshop to the forth International Conference on Language Resources and Evaluation (LREC'04)*, 2004, pp. 1–6.

[7] Á. Herrero Blanco, *Escritura alfabética de la lengua de signos española: once lecciones*. Publicaciones de la Universidad de Alicante, 2003.

[8] W. C. Stokoe, "Sign Language Structure: An Outline of the Visual Communication Systems of the American Deaf," *Studies in linguistics: Occasional papers*, vol. 8, 1960.

[9] V. Sutton and A. Frost, *SignWriting: sign languages are written languages!* Center for Sutton Movement Writing, 2008.

[10] S. M. Thiessen, "A Grammar of SignWriting," M.S. thesis, University of North Dakota, May 2011.

[11] A. C. da Rocha Costa and G. P. Dimuro, "SignWriting-based sign language processing," in *Gesture and Sign Language in Human-Computer Interaction*, 2002, pp. 202–205.

[12] E. Verdu Perez, B. C. Pelayo García-Bustelo, M. Á. Martínez Sánchez, R. González Crespo, et al., "A system to generate signwriting for video tracks enhancing accessibility of deaf people," *International Journal of Interactive Multimedia and Artificial Intelligence*, 4 (6), pp. 109–115, 2017.

[13] O. Koller, H. Ney, and R. Bowden, "May the force be with you: Force-aligned SignWriting for automatic subunit annotation of corpora," in *2013 10th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, 2013. DOI: 10.1109/fg.2013.6553777.

[14] C. Branchini and L. Mantovan, *A Grammar of Italian Sign Language (LIS)*. Fondazione Università Ca' Foscari, Dec. 2020. DOI: 10.30687/978-88-6969-474-5.

[15] Á. Herrero Blanco, *Gramática didáctica de la lengua de signos española (LSE)*. Ediciones SM, 2009.

[16] J. Langer, J. Andres, M. Benešová, and D. Faltýnek, *Quantitative linguistic analysis of Czech sign language*, 1st ed. Univerzita Palackého v Olomouci, 2020. DOI: 10.5507/pdf.20.24457277.

[17] S. K. Liddell and R. E. Johnson, "American Sign Language: The Phonological Base," *Sign Language Studies*, vol. 1064, no. 1, pp. 195–277, 1989.

[18] P. Eccarius and D. Brentari, "Handshape coding made easier: A theoretically based notation for phonological transcription," *Sign Language & Linguistics*, vol. 11, pp. 69–101, 2008.

[19] A. F. G. Sevilla and J. M. Lahoz-Bengochea, "A different description of orientation in sign languages," *Procesamiento del Lenguaje Natural*, 62, pp. 53–60, 2019.

[20] V. Sutton, *Lessons in sign writing: Textbook*, Fourth. Center for Sutton Movement Writing, 2014. [Online]. Available: <https://www.signwriting.org/archive/docs2/sw0116-Lessons-SignWriting.pdf>.

[21] R. W. Smith, "History of the Tesseract OCR engine: What worked and what didn't," in *Document Recognition and Retrieval XX*, vol. 8658, Spie, Feb. 2013, pp. 1–12. DOI: 10.1117/12.2010051.

[22] S. Amraee, M. Chinipardaz, M. Charoosaei, and M. A. Mirzaei, "Handwritten Logic Circuits Analysis Using the YOLO Network and a New Boundary Tracking Algorithm," *IEEE Access*, vol. 10, pp. 76 095–76 104, 2022. DOI: 10.1109/access.2022.3192467.

[23] A. F. G. Sevilla, A. D. Esteban, and J. M. Lahoz-Bengochea, "Quevedo: Annotation and processing of graphical languages," in *Proceedings of the Workshop on Representation and Processing of Sign Language, Workshop to the forth International Conference on Language Resources and Evaluation (LREC'22)*, 2022.

[24] G. H. Granlund and H. Knutsson, *Signal processing for computer vision*. Kluwer Academic Publishers, 1995.

[25] N. O'Mahony, S. Campbell, A. Carvalho, et al., "Deep Learning vs. Traditional Computer Vision," in *Advances in Computer Vision*, K. Arai and S. Kapoor, Eds., Springer International Publishing, 2020, pp. 128–144. DOI: 10.1007/978-3-030-17795-9_10.

[26] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," *Communications of the ACM*, vol. 60, no. 6, pp. 84–90, 2017.

[27] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection," 2016, pp. 779–788. (visited on 06/28/2021).

[28] N. Sambasivan, S. Kapania, H. Highfill, D. Akrong, P. K. Paritosh, and L. M. Aroyo, "everyone wants to do the model work, not the data work": Data Cascades in High-Stakes AI," in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 2021, pp. 1–15.

[29] J. Redmon and A. Farhadi, *Yolov3: An incremental improvement*, 2018. arXiv: 1804.02767.

[30] A. F. G. Sevilla, J. M. Lahoz-Bengochea, and A. D. Esteban, *Visse corpus of spanish signwriting*, version 2.0.0, Zenodo, Mar. 2022. DOI: 10.5281/zenodo.6337885. [Online]. Available: <https://doi.org/10.5281/zenodo.6337885>.

[31] R. Kupriev, D. Petrov, S. Pachhai, et al., *DVC: Data Version Control - Git for Data & Models*, Jun. 2021. DOI: 10.5281/zenodo.5037865. [Online]. Available: <https://zenodo.org/record/5037865> (visited on 06/29/2021).



ANTONIO F. G. SEVILLA majored in Mathematics and in Computer Engineering at Universidad Autónoma de Madrid, and went on to study Computational Linguistics at Charles University in Prague and in Malta University.

Currently at Universidad Complutense de Madrid, his main interest lies in the intersection between knowledge-based and empirical approaches in AI, and specially in Computational Linguistics.

His current research focus is on Sign Languages, their representation and their processing. He is the principal investigator of the project "Visualizing SignWriting" aiming to produce tools for improving the inclusion of SignWriting in the digital world.

He likes to combine scientific research with the development of computational solutions, some of which can be found online as open-source.



ALBERTO DÍAZ ESTEBAN works as Associate Professor (Profesor Titular de Universidad) at the Universidad Complutense de Madrid. He is also a member of the Instituto de Tecnología del Conocimiento.

His research is oriented to natural language processing, in particular to automatic summarization, automatic description of images, information extraction from medical reports and normative documents, and detection of interactions between drugs from text. He is also interested in the automatic processing of Sign Languages, in particular, on SignWriting.



JOSÉ MARÍA LAHOZ-BENGOECHEA is Associate Professor of Spanish Linguistics at the Complutense University of Madrid. His research focuses on different aspects of phonetics and phonology of both oral and sign languages. His dissertation was awarded best dissertation of the year by the Spanish Society of Linguistics.

In 2021 he was awarded a Leonardo Grant for Researchers and Cultural Creators by the BBVA Foundation and thanks to that he is PI of the project “SSL Signary: A parametric dictionary of Spanish Sign Language”.

He has been a visiting scholar in Centre National de la Recherche Scientifique, Paris; University of California, Berkeley; University of California, Los Angeles; MARCS Institute for Brain, Behaviour and Development at Western Sydney University; and Universidad de Navarra.

He is the founder and director of GRIFFOS (Research Group on Phonetics and Phonology of Oral and Sign languages).

...