

The Synthesis of Complex Shape Deployments in Sign Language

Michael Filhol^{1,2}, John McDonald³

¹ Université Paris–Saclay, France

² CNRS, LIMSI, Orsay, France

³ DePaul University, Chicago, IL, USA

michael.filhol@limsi.fr, jmcdonald@cdm.depaul.edu

Abstract

Proform constructs such as classifier predicates and size and shape specifiers are essential elements of Sign Language, but have remained a challenge for synthesis due to their highly variable nature. In contrast to frozen signs, which may be pre-animated or recorded, their variability necessitates a new approach both to their linguistic description and to their synthesis in animation. Though the specification and animation of classifier predicates was covered in previous works, size and shape specifiers have to this date remained unaddressed. This paper presents an efficient method for linguistically describing such specifiers using a small number of rules that cover a large range of possible constructs. It continues to show that with a small number of services in a signing avatar, these descriptions can be synthesized in a natural way that captures the essential gestural actions while also including the subtleties of human motion that make the signing legible.

Keywords: Sign Language Synthesis, Proforms, Classifiers, Size and Shape Specifiers, Avatars

1. Introduction

One of the unique aspects of Sign Language (SL) is its ability to make use of the signing space to locate, link and depict discourse entities in a dynamic manner. They involve iconic projections of topological relationships in the signing space, and symbolic use of spatial anchors for semantic relationships between them (Liddell, 2003; Johnston and Schembri, 2007).

The literature on such productive geometric (spatial) constructions often distinguishes at least two major types (Schembri, 2003; Woll, 2007; Zwitserlood, 2012):

- **classifier predicates**, involving language-specific handshapes (classifiers) that can be placed or moved to show relative positions and movements of semantically typed discourse entities (a person placed here, a car moving this way...);
- **size and shape specification**, involving language-constrained handshapes (SASSes) to describe the shape of an object and deploy lines or surfaces in space (e.g. the neck and body of a large vase).

They all have one purpose in common: description of relationships geometrically projected in the signing space, using handshapes conveying some semantic classification of what they stand for (person, vehicle, flat surface, small flat round object...). More than a handshape, they sometimes bring the whole arm into play (e.g. placing a tree), or a pair of handshapes for a single large object (e.g. a frame on wall). To avoid confusion in this paper, we will call all of these instances, whether used as a classifier or a SASS, a “**proform**”. Proforms are mostly chosen from a language-specific list, e.g. (Vicars, 2020) for ASL, but one can observe others created on the fly.

In previous work, we covered a first set of SL constructions involving proforms (Filhol and McDonald, 2018; McDonald and Filhol, 2019) to demonstrate:

- the powerful geometric abstraction potential of AZee,

a Sign Language modelling approach and description language;

- and the multi-track ability of Paula, a sign synthesis and animation rendering system.

Both designed to allow for parallel tracks and control of all body articulators, they have proven to be well-suited for one another (Filhol et al., 2017).

Among other concepts, we defined a *place-classifier* rule producing a small downward “settling” movement, whose meaning is to anchor an entity at a chosen point in the signing space, and a *move-classifier* rule producing a movement along a path in space, whose meaning is the displacement of the represented entity. This work mostly fell into the first type of proform use mentioned above: classifier predicates. The work reported in sections 2. and 3. of this paper addresses typical constructs of the second type, i.e. shape deployments, where proform movements outline shapes without meaning displacement. Then, in sections 4. through 6., it introduces new avatar animation techniques that facilitate the synthesis of these deployments while avoiding the robotic motion and unnatural postures seen in previous avatar synthesis systems.

2. Simple deployments

First, let us look at various examples of shape deployments¹. Video “curtains” depicts two striped curtains hanging down from above the signer’s face. There are two similar instances of deployments in the video, captured and annotated in figure 1. The first one delimits two sections alongside a window where the curtains are located, with a proform we shall call *thickness-medium*, useful to deploy strips of medium-sized breadth longitudinally. The second one draws stripes on each of these sections, with an eponymous proform *parallel-lines* that extends and spreads the fingers.

¹Videos for quoted example names in this paper are available at <https://doi.org/10.5281/zenodo.3708057>.



Figure 1: Two shape deployments in video “curtains”.



Figure 2: Two shape deployments in video “table”.

Video “table” starts by spreading a table surface with two flat hands (proform *flat-surface*) moving outwards from a point in the middle. This is immediately followed by a second deployment of the same table with a similar movement, only with a proform change to *thickness-medium*, which gives a second point of view on the table, this time as a delimited oblong shape. Annotated still images are given in figure 2.

Video “cupboard” is from the description of a piece of furniture against a wall. Its back is represented by the flat weak hand in the background (proform *flat-surface*), and a glass window in the front deployed on the strong hand in a vertical plane in the foreground (same proform). An annotated still shot of the relevant deployment construction is given in figure 3.

Continuing our prior work of formalising rules in AZee to represent more of the possible productions, we applied the AZee search methodology to a corpus containing many such examples (Benchiheub et al., 2016). This methodology is based on alternating searches for articulated forms and interpreted meanings in a corpus, and retrieving stable meaning–form associations. Branching and inverting each iteration by carrying over the common form/meaning counterpart as the starting point of the next, we establish production rules usable for synthesis (Hadjadj et al., 2018). Each production rule is a function that determines the forms to articulate for an identified meaning when applied (and given its arguments if any).




Figure 3: Shape deployment in video “cupboard”.

Looking at our examples, this method would lead to:

- interpret every annotated path in figures 1, 2 and 3 as **meaning** the depicted path (say P) deployed by the articulated proform (say prf);
- observe the same **form** (except for the differences accounted for by P and prf) for every instance, i.e. prf follows P with invariable dynamics, and the eye gaze is directed to the position of prf .

This observation warrants the definition of an AZee production rule, which we shall name *deploy-shape*, function of P and prf , carrying the meaning identified above and producing the form described above.

Say we now define a path $P_{L \rightarrow R}$ from left to right in the signing space. Applying the rule *deploy-shape* to it with, say, proform *flat-surface* as exemplified in E1 below would generate a sign score specifying a horizontal left-to-right sweep of the proform, then to be interpreted—although out of context here—as a flat surface deployed along the given path.

E1 *deploy-shape*($P_{L \rightarrow R}$, *flat-surface*)

All deployments in the proposed video example list apply *deploy-shape* on some level, each with its own argument values.

Examples “curtains” and “table” involve moments when two paths are deployed at the same time. Our earlier work introduced rule *simultaneous* to place two cups symmetrically on the table, or a knife and a fork on either side of a plate (McDonald and Filhol, 2019). Given two statements, the interpretation (meaning) of *simultaneous* is that both of them are true or happen at the same time, and its production (form) is the articulation of both simultaneously—typically on either side of the body. Shapes depicted by articulating two simultaneous deployments like in videos “curtains” and “table” verify the form produced by *simultaneous*, but also the corresponding interpretation: both parts of the shape exist at the same time, whether they are part of a same, physically continuous object (e.g. “table”) or not (e.g. “curtains”). The rule *simultaneous* is therefore well suited to capture this type of production, applying a *deploy-shape* on each hand.

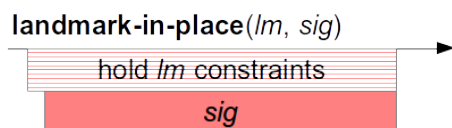


Figure 4: Box diagram for production rule *landmark-in-place*.

For example, given the two manual paths P_a and P_b shown in figure 1, the corresponding parallel deployments in video “curtains” can be represented by the AZee expression E2 below, with $prf=thickness-medium$ for the first instance, and $prf=parallel-lines$ for the second². AZee terms “*s*” and “*w*” stand for strong and weak side respectively.

E2 *simultaneous*(*deploy-shape*(P_a , $prf(hand=s)$), *deploy-shape*(P_b , $prf(hand=w)$))

In video “cupboard”, a hand is fixed, next to which a shape is deployed with the other hand. Sometimes, such examples make the fixed hand one side or end of the deployed shape, while the other takes care of the deployment from that point. Video “bedroom-walls” is an example of this, which we deal further with in the next section. In all of these instances, the fixed hand is interpreted as an active landmark in space, relative to which the rest is signed and potentially located as long as it is held in place. To capture this meaning–form association, we propose the new AZee production rule named *landmark-in-place*, function of a postural constraint lm for the held landmark and a signed piece of discourse sig , to be interpreted in the spatial context activated by lm . Its produced form is that of sig with lm installed just before sig starts, and held until sig ends (see fig. 4).

Given the location K of the background landmark and the path P_c shown in fig. 3, video “cupboard” can be accounted for by E3 below.

E3 *landmark-in-place*(*flat-surface*($hand=w$, $loc=K$), *deploy-shape*(P_c , *flat-surface*($hand=s$)))

3. Complex deployments

In this section, we take the challenge one level up and look at more complex shape deployments such as the one in video “bedroom-walls” (fig. 5). It describes the shape of a room by depicting two opposite walls, as illustrated in the diagram in fig. 6. The first wall (red) is made of two sections in a straight angle, while the second (blue) is made of one, reaching further out in distance from the front-most point. The video exhibits three horizontal manual strokes, one for each wall section (AB , BC and DE), plus an intermediate transitioning movement from C to D .

One immediately recognises shape deployments similar to those produced by *deploy-shape* defined above, and the use of *landmark-in-place* overarching the whole description in this case. To this extent the construct is comparable to E3, but the full utterance cannot be captured with a single followed path like in E3 because:

²Note that orientation issues are not specified in E2, for simplification. They will however be dealt with further down.



Figure 5: Shape deployment in video “bedroom-walls”.

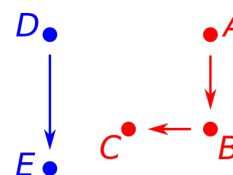


Figure 6: Layout of the wall sections in “bedroom-walls”.

- the manual repositioning from C to D cannot be interpreted as a part of the depicted wall;
- the first two hand strokes do not follow a one-stroke dynamic that could be the result of a single followed path;
- the proform’s orientation changes twice (at B and somewhere between C and D), without the rotation itself following a controlled path curve.

The utterance instead contains a combination of three straight deployments. Each one is describable by an expression similar to E1, but their combination remains an issue. They come one after the other, but following the AZee principle, no signed sequence is produced without identifiable production rules to justify it. In other words, the observed path concatenation must itself be the result of an application of defined production rules.

If we look at the signed forms visible in the video in more detail, we notice a few more clues:

- the first two strokes AB and BC are performed back to back, with shorter durations than the last or than those in the other videos;
- the last separate stroke DE is produced after a brief hold of the preceding posture (right hand at position C) and a quick but noticeable blink;
- another short hold is visible at position E , also with a blink.

These clues remind us of the forms produced by the rules *each-of* and *all-of*, also introduced in our work on classifiers (Filhol and McDonald, 2018). The rule *each-of* articulates an argument list of items in sequence, holding the final posture of each one for a short moment and appending a quick blink of the eyelids at the end of each hold. Its

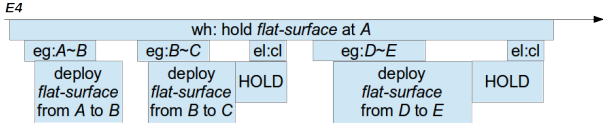


Figure 7: Sign score resulting from AZee expression E4.

meaning is that every element listed is equally true and focused, i.e. each is important, but none more salient than the other. It is a way in SL of building an exhaustive list of separate items, events or clauses of equal status. For example, our table scene description used it to lay out elements placed on a table.

The similar rule *all-of* also articulates an argument list of items in sequence, but each of them with accelerated dynamics (produced faster), and with no held feature in between. Its meaning also involves grouping the listed elements in a set of joint items or clauses, but this time focuses on the set as a whole, removing the relevance from its contained parts. It appeared in the table scene too, for example grouping 4 classifier placements to mean “[set of] 4 plates”, and grouping signs KNIFE and FORK to mean “knife & fork pair”—cutlery items were then positioned pair by pair using two simultaneous classifiers for each pair.

By combining these two rules with *deploy-shape* and *landmark-in-place*, it is now possible to capture the form of “bedroom-walls” exactly, with expression E4 below³. The colours correspond to those in fig. 6, and are applied to the parts describing the matching wall sections.

E4 *landmark-in-place*(standing-wall(*nrm=lat*, hand=*w*, loc=*A*), each-of(*all-of*(*deploy-shape*($P_{A \rightarrow B}$, *standing-wall*(*nrm=-lat*)), *deploy-shape*($P_{B \rightarrow C}$, *standing-wall*(*nrm=fwd*))), *deploy-shape*($P_{D \rightarrow E}$, *standing-wall*(*nrm=-lat*))))

In this expression:

- *standing-wall* is a short-cut to specify *flat-surface* with the fingers up, leaving open the horizontal vector *nrm* which defines the plane in which the wall lies (normal orientation);
- $P_{X \rightarrow Y}$ is the straight path from point *X* to *Y*;
- *lat* an ipsilateral vector (pointing to the strong side), *fwd* a vector pointing forward (outwards from the body).

Expression E4 evaluates to the sign score (time line of signed forms) represented in figure 7, which matches the articulations visible in video “bedroom-walls” well. Moreover, like any AZee expression combining production rules, E4 not only produces forms but also conveys a composite meaning resulting from the semantic combination of the rules nested in the expression. In this case, E4 can be broken down into the following interpretation:

Drawing the scene from corner A of the room [landmark-in-place], [there are] two separate wall sections [each-of]: one from A, made of two subsections [all-of] AB and BC, and the other from D to E.

It appears that E4 gives the whole construct an interpretation that is entirely compatible with the meaning of the video.

4. Animating AZee

Given the above linguistic representations for sequences of deployments, the next task is to synthesize it as animation on a human avatar. The data provided by AZee consists of a series of timed blocks as in fig. 7, which are organized hierarchically. Each of these blocks controls different processes on the avatar’s anatomy, which may all affect overlapping parts of the avatar’s anatomy. For instance, the specifications for both the strong and weak hand “deploy flat surface” processes will affect the following parts of the anatomy (McDonald et al., 2017):

- the hands and arms to deploy the shape;
- the neck and eyes to direct gaze to the shape;
- the torso to support both of these processes.

The avatar must not only be able to schedule the sequence of required postures, it must be able to combine and blend their effects on the anatomy seamlessly.

More importantly, while each of these blocks contains animation information (e.g. as a sparse set of key body postures, or a mathematical procedure), this information is necessarily an abstraction which defines just enough to carry the meaning, but which leaves out details of human motion that are essential to making signing look natural, legible and more human. This tug-of-war between the sparseness of the linguistic abstraction and the richness of human motion, has long plagued efforts to build avatars that synthesize sign directly from linguistic descriptions. However, the essential nature of this interplay has made it a key element of the bridge we have built between linguistics and animation, and has led to the present study.

Recall that the goal of the Paula sign synthesis system has been to leverage two key elements in an effort to animate sign legibly and naturally:

- the structure from a linguistic description of sign (Wolfe et al., 2015);
- the experienced eye and hand of an artist (Wolfe et al., 2011).

The previously proposed bridge (Filhol et al., 2017) from AZee sign descriptions to the animation of the Paula avatar is built on these two pillars through a system of templated shortcuts. These allow Paula to construct animations from larger blocks of motion, rather than the individual posture specifications that have driven prior avatars. These motion

³Unlike in previous examples, E4 includes orientation specifications, which change over the course of the depictions.

blocks can be of a range of types including mathematical procedures, pre-animated sequences, and hybrid procedures that draw significant posture and motion data from a small number of pre-animated sources.

Extending this bridge to animate the shape deployments described above becomes clear by reviewing the methods that the bridge has used for various types of discourse. When animating a frozen sign (whose form rarely changes other than for co-articulation or for ease of production) the system is free to shortcut to a pre-animated gestural unit that can be dropped in place of the AZee block and can then be blended with other elements.

These shortcuts were expanded to include templatable information in the prior case of classifier placement and movement (Filhol and McDonald, 2018). Here, the position of the placement and the direction of the movement change from production to production. However there are many other parameters that are left unconstrained by the linguistic description, and so may be set to whatever the animation system deems appropriate. This allowed Paula to leverage configuration data from an artist generated pose for the proform, which could include:

- the configuration of the hand for the proform;
- an orientation for the hand;
- the natural configuration of the shoulder and elbow;
- the accompanying configuration of the torso that supports the pose.

In this simpler situation, in comparison to the present study, the system was able to leverage this artist data because it is left unconstrained by AZee, i.e. not specified linguistically. When a parameter on the avatar is unconstrained, the system is free to choose a value for a parameter such as the height of the elbow, and wrist orientation that is comfortable and natural. It is precisely this comfort and naturalness provided by the artist that is one of the strengths of the Paula avatar. The additional fact that there are a limited number of commonly used classifiers, and the fact that the generic pose need only be set up once, makes this possible and not an undue burden on the artist.

Prior systems for generating sign movements directly from linguistic descriptions such as HamNoSys (Hanke, 2004) relied on automatic computations from inverse kinematics solutions, techniques originally designed to control robots (Buss, 2004), which contributed to the robotic nature of avatar motion from pure synthesis.

By leveraging an artist's eye for the pose and motion of the human body, the AZee-to-Paula bridge has been able to produce far more natural motion than prior systems were capable of. Another factor that contributes to the naturalness of motion generated by this system is the collection of ease controls for smooth motion control that are exploited as in (McDonald and Filhol, 2019).

5. Animating Deployments

The present study centers on a collection of shape deployments which constrain the system to a far greater degree than the prior classifier movements. These include:



Figure 8: Extreme wrist rotations in avatars

1. placement of objects, constrained in orientation that deviate significantly from the artist pose;
2. the deployment of surfaces in space described above, which will follow complicated orientations as the hand traces the surface shape.

Both of these situations will fully constrain the orientation of the hand in space relative to the body. For example, consider the deployment of a wall situated in front of the signer, and extending from left to right. The signer's palm will face the wall, i.e. out from the body, with fingers pointed up to show the wall's surface. The hand will then move toward the right to show the extension of the wall as laid out in figure 6, while maintaining the orientation.

These added constraints may at first seem like an advantage for the avatar, since it has far fewer unconstrained elements to fill. Unfortunately, from the perspective of producing natural motion it actually puts a straight-jacket on the avatar, forcing it into unnatural postures because of the coarseness of the linguistic specification as seen in the examples in figure 8 which are examples from the avatars described in (Kipp et al., 2011) and (Elliott et al., 2008). Such postures can even plague motion-capture derived signing due to the need for retargeting which often relies on the same kind of inverse kinematic solutions (Awad et al., 2009).

Paula would encounter the same issue in the second segment of expression E4, where the linguistic description specifies that the strong hand begin its motion at a medium distance in front of the weak shoulder with the hand facing forward and up against the horizontal wall. If the animation system were to attempt to orient the hand perfectly along these cardinal axes, the result would be seen in the left image in figure 9.

Of course, the human body never positions itself with such precision and considerations of comfort and strain will modify both the desired position and orientation. Notice that in figure 5, the hand is not pointed straight up nor is it facing perfectly forward. Yet, both the linguistic abstraction and the way that the resulting position is perceived by a viewer is consistent with an upward pointing-forward facing hand. The right image in figure 9 illustrates a more natural relaxed configuration of the hand.

The Paula system has, to this point, avoided this problem



Figure 9: Literal vs. relaxed interpretation of E4

through its reliance on an artist’s touch. Unfortunately, in the present application, the infinite collection of possible positions and orientations precludes using an artist defined pose. So, the system must fall back on a more traditional application of inverse kinematics that treats the hand as a block to be positioned and oriented in space, with the shoulder and elbow rotated to place it there. The wrist is then forced into the orientation needed to obey the constraint, even if that orientation would break a human wrist.

6. Relaxing wrist orientations

To avoid such unnatural wrist postures, the new system introduces a relaxation algorithm which balances the linguistically specified spatial orientation with the perceived strain that the hand would be under. The human wrist’s comfortable range of motion is a good start for this and is roughly the following (Gates et al., 2016).

- -40° to 40° for wrist flexion/extension;
- -25° to 25° for wrist ulnar/radial deviation;
- -60° to 60° for wrist flexion/extension.

To relax the hand and allow a more comfortable posture modification the Paula system applies a penalty to wrist angles outside the comfort range of motion⁴. Outside of this range, the angle will increase at a slower rate than would be specified linguistically until it reaches the maximum rotation of the joint, see figure 10.

If $\pm v_0$ is the discomfort free range, $\pm v_1$ is the maximum range of a joint, and $v = \frac{2(v_1 - v_0)}{\pi}$ then this can be achieved with a relaxation of:

$$v_0 + V \cdot \arctan\left(\frac{x - v_0}{V}\right)$$

Applying such a relaxation penalty to each of the rotations at the wrist yields the more natural pose seen on the right of fig. 9. The hand still reads as facing the far wall, but the wrist is no longer strained.

⁴In our implementation, these angles are set somewhat smaller than the physical ranges ($\pm 30^\circ$, $\pm 15^\circ$, $\pm 50^\circ$ respectively) since the skin of the avatar shows strain slightly earlier than a human wrist would.

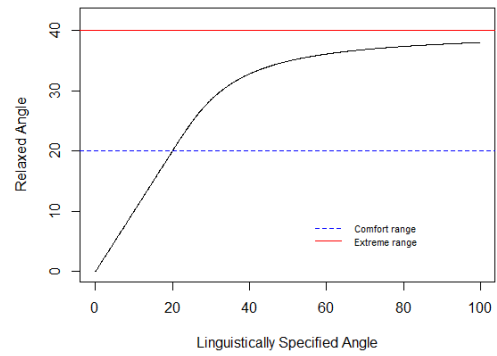


Figure 10: Angle Relaxation Function.




Figure 11: Synthesis of E2 with the Paula system.

7. Results

After implementing the features explained above on the Paula side of the system, the AZee parser was run on expressions E2 (*parallel-lines* in “curtains”), E3 (“cupboard”) and E4 (“bedroom-walls”). The respective animations obtained can be found at <https://doi.org/10.5281/zenodo.3708057>, and illustrated in figures 11, 12 and 13.

The animations in these examples show that the avatar is able to follow the paths specified by the linguistic descrip-



Figure 12: Synthesis of E3 with the Paula system.



Figure 13: Synthesis of E4 with the Paula system.




Figure 14: Repeated classifier placement with a landmark “in place”.

tion, and further that the relaxation at work in these examples provides a naturalness that has previously not been achieved when driven directly from a linguistic description. Aside from the practical achievement of enlarging the set of AZee expressions that Paula is able to generate, we highlight that the type of structures rendered here, namely statements involving proforms, are of a non-fixed geometric kind that no other SL synthesis system has yet covered. By using geometric constructions as arguments of production rules, i.e. points, vectors... and transformations like scaling or translating, one can write an infinite number of AZee expressions, semantically composed and accounting for the ability in SL to make a productive, on-the-fly use of signing space.

Equally important is the fact that this expressive power emerges from a very small set of production rules. Indeed, the list of rules appearing in the reported AZee expressions in all of our works on proforms combined, in addition to the proforms themselves, are:

- *place-classifier*, producing the small “settle” movement ending at the point where to place the proform;
- *move-classifier*, making a proform follow a path and meaning to depict the displacement;
- *deploy-shape*, making a proform follow a path and meaning to depict the drawn shape;
- *simultaneous*, producing two statements at once and meaning that they are simultaneous;
- *landmark-in-place*, producing a statement while a fixed landmark is active;
- *each-of*, producing a list of separate items, each with equal importance;
- *all-of*, producing a set of items, with focus on the formed set.

In other words, the built-in geometric operators, a few proforms and seven production rules were enough to cover a large array of proform placements and shape deployments. This set includes productions whether they are used by themselves or in relation to one another, and whether they consist of a single stroke or of multiple paths.



Figure 15: Frozen sign RUSSIA while laying out a map with a landmark “in place”.

Besides, in the AZee paradigm, whenever both a form and a meaning is found in a corpus to match those identified for a defined production rule, one can label the utterance as an application of the rule, provided its arguments can be identified as well. For example, we know that *landmark-in-place* can combine a given set of articulatory constraints (landmark argument *lm*) with any signing score (argument *sig*), which is to be interpreted in the spatial context of *lm*. We have seen this used to locate SASS deployments (e.g. “cupboard”), possibly repeated (e.g. “bedroom-walls”), but it can also be found with classifier placements or movements, also possibly repeated as exhibited in video “wine-bottles” (see fig. 14). What is more, it can combine with more complex scores mixing even dictionary signs like RUSSIA in example “map-layout”, shown in figure 15.

Therefore, *landmark-in-place*, originally created for proforms in this paper, is not limited to proform constructions, let alone only to one type. Instead, it is much more generally applicable and transparently encompasses features that traditionally called for new concepts, like “buoys” (Liddell, 2003). Plus, *each-of*, *all-of* and *simultaneous* were created for expressions without proforms, and now used in this context.

In the light of this, we wish to emphasise the benefit of the general approach. Breaking down structures to arbitrarily deep levels and factoring elements into production

rules whenever consistent form–meaning associations are observed can provide insight on traditional linguistic categories.

8. Conclusion

This paper set out to extend the AZee coverage of Sign Language constructions depicting shapes deployed in space, as well as their animation with the Paula avatar. On the linguistic side, we introduced new rules such as *deploy-shape* and *landmark-in-place*, and reused prior rules like *simultaneous* when they fit the observed forms and carried the right meaning. On the synthesis side, we implemented new features such as geometric orientation of proforms and wrist relaxation to add naturalness to the postures where abstract linguistic specifications would otherwise lead to robotic or unnatural positions. With these efforts we managed to enlarge the set of proform constructions accounted for, from both of the AZee linguistic model and the Paula synthesis perspectives.

The naturalness in the output animation and expressive power in the input representation is encouraging and serves as an important validation step of both the linguistic model and the animation engine. Further, the synergy in the overall system drives forward the state of the art for both animation synthesis and linguistic representation, expanding the ability of avatars to produce even generative sign constructs such as proforms directly from linguistic descriptions.

Despite these significant gains in coverage, some aspects of proform constructions are still missing in our system, like proforms following curves, e.g. depicting a car taking a curve in the road. Future work will address these following a similar incremental methodology.

Bibliographical References

- Awad, C., Courty, N., Duarte, K., Le Naour, T., and Gibet, S. (2009). A combined semantic and motion capture database for real-time sign language synthesis. In *International Workshop on Intelligent Virtual Agents*, pages 432–438. Springer.
- Benchiheb, M., Berret, B., and Braffort, A. (2016). Collecting and analysing a motion-capture corpus of french sign language. In *7th International Conference on Language Resources and Evaluation (LREC), Workshop on the Representation and Processing of Sign Languages*, Portoroz, Slovenia.
- Buss, S. R. (2004). Introduction to inverse kinematics with jacobian transpose, pseudoinverse and damped least squares methods. *IEEE Journal of Robotics and Automation*, 17(1-19):16.
- Elliott, R., Glauert, J. R., Kennaway, J., Marshall, I., and Safar, E. (2008). Linguistic modelling and language-processing technologies for avatar-based sign language presentation. *Universal Access in the Information Society*, 6(4):375–391.
- Filhol, M. and McDonald, J. (2018). Extending the AZee-Paula shortcuts to enable natural proform synthesis. In *Workshop on the Representation and Processing of Sign Languages*, Miyazaki, Japan, May.
- Filhol, M., McDonald, J., and Wolfe, R. (2017). Synthesizing sign language by connecting linguistically structured descriptions to a multi-track animation system. *Universal Access in Human-Computer Interaction, Lecture Notes in Computer Science (Springer)*, 10278:27–40.
- Gates, D. H., Walters, L. S., Cowley, J., Wilken, J. M., and Resnik, L. (2016). Range of motion requirements for upper-limb activities of daily living. *American Journal of Occupational Therapy*, 70(1):7001350010p1–7001350010p10.
- Hadjadj, M., Filhol, M., and Braffort, A. (2018). Modeling French Sign Language: a proposal for a semantically compositional system. In ELRA, editor, *International Conference on Language Resources and Evaluation*, Miyazaki, Japan, May. ELRA.
- Hanke, T. (2004). Hamnosys-representing sign language data in language resources and language processing contexts. In *LREC*, volume 4, pages 1–6.
- Johnston, T. and Schembri, A. (2007). *Australian Sign Language (Auslan): an introduction to sign language linguistics*, volume 117. Cambridge, July.
- Kipp, M., Heloir, A., and Nguyen, Q. (2011). Sign language avatars: Animation and comprehensibility. In *International Workshop on Intelligent Virtual Agents*, pages 113–126. Springer.
- Liddell, S. (2003). *Grammar, gesture and meaning in American Sign Language*. Cambridge University Press.
- McDonald, J. and Filhol, M. (2019). Fine Tuning Dynamics in Contextualized Proform Constructs from Linguistic Descriptions. In *International Workshop on Sign Language Translation and Avatar Technology*, Hamburg, Germany, September.
- McDonald, J., Wolfe, R., Johnson, S., Baowidan, S., Moncrief, R., and Guo, N. (2017). An improved framework for layering linguistic processes in sign language generation: Why there should never be a “brows” tier. In *International Conference on Universal Access in Human-Computer Interaction*, pages 41–54. Springer.
- Schembri, A., (2003). *Perspectives on Classifier Constructions in Sign Languages*, chapter Rethinking ‘classifiers’ in signed languages, pages 3–34. Psychology Press.
- Vicars, W. (2020). ASL classifiers: <https://www.lifeprint.com/asl101/pages-signs/classifiers/classifiers-frame.htm>.
- Wolfe, R., McDonald, J., and Schnepf, J. C. (2011). Avatar to depict sign language: Building from reusable hand animation. In *International Workshop on Sign Language Translation and Avatar Technology*, Berlin, Germany.
- Wolfe, R., McDonald, J., Moncrief, R., Baowidan, S., and Stumbo, M. (2015). Inferring biomechanical kinematics from linguistic data: A case study for role shift. In *Symposium on Sign Language Translation and Avatar Technology (SLTAT)*, Paris, France.
- Woll, B. (2007). The linguistics of sign language classifiers: phonology, morpho-syntax, semantics and discourse. *Lingua: International Review of General Linguistics*, 117(7):1159–1353, July.
- Zwitserslood, I., (2012). *Sign Language: an International Handbook*, chapter Classifiers, pages 158–186. Mouton de Gruyter, Berlin.