

Extending the AZee-Paula Shortcuts to Enable Natural Proform Synthesis

Michael Filhol^{1,2}, John McDonald³

¹ LIMSI (CNRS), ² Université Paris–Saclay, ³ DePaul University

¹ Univ. Paris–sud, bât. 508, 91400 Orsay, France

³ 243 S. Wabash, Chicago, IL, USA

michael.filhol@limsi.fr, ³ jmcDonald@cdm.depaul.edu

Abstract

Proform structures such as classifier predicates have traditionally challenged Sign Language (SL) synthesis systems, particularly in respect to the production of smooth natural motion. To address this issue a synthesizer must necessarily leverage a structured linguistic model for such constructs to specify the linguistic constraints, and also an animation system that is able to provide natural avatar motion within the confines of those constraints. The proposed system bridges two existing technologies, taking advantage of the ability of AZee to encode both the form and functional linguistic aspects of the proform movements and on the Paula avatar system to provide convincing human motion. The system extends a previous principle that more natural motion arises from leveraging knowledge of larger structures in the linguistic description.

Keywords: Sign Language Linguistics, Proforms, Classifiers, Synthesis, Avatars

1. Introduction

Producing natural synthesis of sign languages using an avatar is a goal with far-reaching applications for Deaf-hearing communication including improving sign language education tools, anonymizing online communication in sign language, and enabling translation for situations where hiring a certified interpreter is impossible. To support all of these applications, a sign synthesizer must be able to express all aspects of sign language including the full range of body signals used to communicate.

Sign languages use a range of linguistic processes to communicate, the most basic of which being those listed in sign dictionaries. These are gestural units that have a standardized and stable meaning–form association. In addition, signers use a range of grammatical processes for rich, natural communication. All of these processes are communicated through signals involving the arms, eyes, face, torso and neck of the signer. Natural sign synthesis remains a challenge partially because of the fact that these structures can overlap and interact on the body (Weast, 2011) and animating such structures requires leveraging both sign language linguistics and knowledge of human motion (Braf-fort et al., 2015).

Most current sign synthesis systems are able to animate a stream of lexical signs (Wolfe et al., 2011; Elliott et al., 2008; Lombardo et al., 2011). However, more freehand constructs in sign languages such as classifier placement and movement or size and shape specifiers remain a challenge because of the inherent variability in form expressed through the signer’s body. These structures use configurations on the signer’s body known as proforms, wherein a part of the signer’s body stands in for an object and is often iconic of its shape. For example, in American Sign Language (ASL) and French Sign Language (LSF), the index finger oriented vertically will represent a standing person, whereas a “C” handshape will represent a cylindrical object such as a glass (Liddell, 2003). Such proforms can be used to express the placement or movement of objects or to

describe their size and shape.

Proform structures have traditionally been a challenge for synthesis systems, both from a linguistic and from an animation standpoint, because they are highly productive. The placement and movement of body articulators cannot easily be captured by pre-set configurations, and so cannot be pre-animated or recorded. Thus systems relying on motion capture (Gibet et al., 2011) or reusable hand animation (Wolfe et al., 2011) must fall back on more primitive synthesis techniques. In addition, while many proforms are predefined for a given sign language, variations in the shape or configuration of an object can be expressed through a near infinite range of hand or body configurations, especially when one considers size and shape specifiers (Liddell, 2003).

Prior efforts have relied either on synthesis from phonetic components (López-Colino and Colás, 2011) or on predefined templates that encapsulate a limited set of standard proforms (Huenerfauth, 2004). In both cases the results were far from natural involving awkward body configurations and robotic motion. The present work addresses three weaknesses that contribute to the robotic nature of the resulting animations.

1. Body gestures are never fully specified by the linguistics since there are a range of body configurations that can satisfy given linguistic constraints (Filhol and Braffort, 2006b). Synthesis systems have had to fill in the missing body constraints for example by overspecifying the motion linguistically.
2. Synthesized motion has been limited to the avatar’s arms, whereas natural human arm motion is always accompanied by supporting torso and clavicle motion. (McDonald et al., 2016a).
3. Motion specifications were limited to key positions for the handshape and didn’t specify dynamic differences such as acceleration or speed profiles that can profoundly affect perception.

To fully specify proform movement and placement, and to produce natural motion, the synthesizer must take into consideration both:

- the linguistically defined constraints, which abstract human motion into meaningful gestural units;
- the range of human motion that accompanies such gestural units, but which are not encoded linguistically.

This paper extends the work in (Filhol et al., 2017) which sought to bridge between a structured linguistic model of sign and a multitrack sign animation system. The system described here achieves natural linguistically driven proform motion through two key features. First, by separating the task into separate linguistic and animation components, it allows the linguistic component to encode the necessary information for the proform while allowing it to remain underspecified at a geometric level. This gives the avatar the needed freedom to move naturally within linguistic constraints. Second, it builds on the prior model’s principle that natural motion is best achieved with large linguistic structures rather than from very basic phonetic specifications. The next section explores the perspectives of each system on proforms and what each offers for building a combined synthesizer.

2. Perspectives of the two systems on proforms

2.1. AZee for descriptions

The Sign Language description model AZee (Filhol et al., 2014) has several advantages that are relevant to the use of proforms in SL. The major one is its fully embedded geometric system that allows to build and describe points, vectors and paths in the signing space. From the origin of its predecessor Zebedee, geometric specification of body locations and skeletal orientations as points and vectors in an affine 3D space¹ have been an essential feature of the descriptions (Filhol and Braffort, 2006a). More than an alternative style of body posture description, such geometric approach accounts for at least three notable features of SL, which are difficult to capture with other, e.g. parametric, models.

First, it does not rely on a discrete set of points for locations in space, or directions for orientations. It allows to define geometric objects in a continuous space. In other words given two points, it is always possible to take the midpoint of the two. When making free productive uses of space, e.g. placing proforms to indicate relative positions, it is therefore possible to account for any relative placement, such as a date “in the middle of” the two boundaries of a delimited period on a time axis.

Second, dependencies between elements of the descriptions are made relevant. This is useful for depicting structures such as the proform placements or movements we are addressing because positions are often relative to (dependant

¹In geometry, an affine space is a vector space with no chosen metric or origin. In our case everything is defined relative to the body, including directions and distances. This allows implementation with any avatar. No body geometry is assumed by AZee.

on) each other. For example, in the predicate “rabbit near and on the right of tree”, the target location point for the rabbit proform is relative to that of the tree. AZee allows to express the rabbit’s position as a geometric translation of the point with the appropriate distance, instead of projecting to a grid defining everything relative to the chest.

Third, not only hands can be specified target locations but any articulator of the body. We have argued the necessity of this in several papers, but it becomes all the more relevant in dealing with depicting structures. Placing two-hand classifiers (e.g. round plates) and placing a full-arm classifier (e.g. tree) pose a problem if we want to consider hand placement alone as changing the classifier inside any expression will require a change of location as well. In AZee, a plate and a tree can be placed at the same location using the same point, even though hands actually end up in completely different locations.

Essentially, AZee is a language to write *rules* mapping invariant and parameterized *forms* to identified semantic *functions*, regardless of the level of linguistic description. Classifier/proform placements are indeed units that are difficult to locate in those terms, as they are arguably both lexical and grammatical constructions, or neither (Johnston categorizes those separately as “partly lexical signs”). AZee bypasses this problem, as any function-to-form link is tackled using the same description model. Like any other rule, a classifier generates a set of articulatory constraints for a consistently interpreted meaning.

For example, the upright index finger shape denotes a standing person, possibly with a wrist orientation depicting the direction in which the person is facing. Because the meaning conveyed with this finger arrangement is consistent, an AZee rule “proform-standing-person” can be defined to specify the appropriate articulatory constraints: index up, possibly facing along a parameterized direction, other fingers closed. As we stated with the AZee approach, only the set of necessary and sufficient constraints are to be specified. In this instance we therefore exclude palm orientation from “proform-standing-person”, because only fingers matter. In contrast, a rule for “tree” would have to include all bones from the tip of the fingers down to the elbow, since not extending and spreading the fingers would result in breaking the meaning.

The AZee approach also encourages and facilitates factoring similar forms into new rules when the interpreted meanings share a common factor across multiple productions. For example, take the placements of entities in the signing space of the kind Liddell glosses with -BE-AT[↓] suffixes. All are produced with a small settling movement towards the surface on which the object is placed (often downwards), and an eye gaze towards the target location. This is true regardless of what the object is, and regardless of what articulator set is conveying the object. In AZee, one would therefore factor this common form, parameterizing the proform *prf* and the location point *loc* to define a rule with semantic function:

“placement of *prf* in space at location *loc*”

producing the form:

“small straight movement of *prf* down to point *loc* + look at *loc* + synchronize eye gaze with a negative time offset”.

Proform *prf* and point *loc* become empty placeholders, and necessary arguments of the created rule.

More factoring can take place, this time with a rule already found and reported elsewhere: the “category” rule, whose semantic function is to give a category in which to interpret a second argument item. It allows to juxtapose “town” and “Berlin”, or “profession” and “bakery” to specify a sort of hypernym for the second item, likely but not necessarily ambiguous on its own. Classifier constructions in LSF often involve juxtapositions of a dictionary sign and a placement like the one described above, applied to a proform. The overall form then, including the juxtaposition, a slight head tilt and a specific inter-sign transition timing, is the same as that specified by the “category” rule. Plus, the meaning is to us all but similar: the first item gives a class of which the second is an instance.

Therefore, from the three rules below:

- category
- proform-vehicle
- place-prf

and a point *loc*, one can build the complex expression:

(E1) category(car(), place-prf(proform-vehicle, *loc*))

producing the sequence traditionally glossed as:

CAR VEHICLE-BE-AT^{loc}

and taking care of the precise timing and adding the gaze towards *loc*.

2.2. Paula for Natural Animations of Sign

The Paula sign synthesis system compliments AZee’s linguistic proposition in supporting natural animations of SL. As described in previous publications, it is a hybrid animation system that supports layering motion from a variety of sources including procedural, keyframe, etc. (McDonald et al., 2017). From the synthesizer’s perspective, proform movements can be modeled as a collection of keyframe data. In this respect, Paula offers a range of features that allow it to produce more natural animation from such data, and allows leveraging animator/sign-expert skills to a larger degree than prior systems. The following features are key to producing natural animations of proform movement from the linguistic specifications:

1. Key postures can be set by using either forward (FK) or inverse-kinematics (IK) systems tuned for sign linguistics (McDonald et al.,). The IK system allows any point defined relative to the hand or arm to be placed at a chosen target in space, or at a chosen site on the avatar’s body, and allows full exploration of the redundant degrees of freedom in the IK chain.
2. Keyframes can be scheduled completely asynchronously on different articulatory chains (McDonald et al., 2017).

3. Interpolations between key postures are accomplished using nonlinear rotation controllers that create transitions following the natural arcs of human motion. These interpolators also allow independent control of speed and trajectory along paths.
4. Procedural techniques automatically allow the torso and shoulders of the system to naturally support and accompany arm movements (McDonald et al., 2016a), and also add sub-linguistic ambient motion (McDonald et al., 2016b).

These features all contribute to allow the synthesizer to produce smooth, natural movement from the linguistic features described above.

The first of these features allows any articulator on or near the arm to be used for targeting, and facilitate the natural positioning of two-handed proforms such as in the *round-plate* classifier example from the last section. The fourth feature in this list extends the avatar’s arm motion through the trunk all the way to the hips, directly addressing the second cause of robotic motion cited in the introduction. Finally and Paula’s nonlinear rotation controllers provide for independent velocity control on articulatory chains, needed to animate the linguistic categories of dynamics, thus addressing the third cause of robotic motion.

To address the underspecified nature of proforms, Paula also offers components that help leverage animator and sign-expert skills as far, and as deep, as possible in the synthesis process. The first component most often used in sign animation is Paula’s Sign Transcriber, which scripts for pre-recordable segments (Wolfe et al., 2011). In spite of the fact that proforms are highly variable, Paula’s Sign Transcriber allows animators to provide significant body posture and movement hints, which the proform generation system can then use to produce more natural animations than would be possible from the linguistic data alone. What remains is to build a coherent bridge between these two systems so that the necessary constraints are communicated while giving the animation system the freedom it needs.

3. Extending “the coarser the better”

3.1. Last proposal (and clarification)

The present work builds on the AZee-Paula bridge proposed in (Filhol et al., 2017), which was based on the principle that working with larger blocks of animation or procedural motion will generate more natural sign synthesis compared to animating from individual constraints and joint settings. The previous bridge mapped from AZee expressions that could be recognized when reading block descriptions output by the AZee parser. This allowed the animation system to shortcut the application of a block if a prerecorded or procedural animation is directly available, rather than reconstructing the entire block’s animation from low-level primitives nested in the generated form description.

This principle, which the authors called “the coarser the better”, relies on AZee’s organization of scores as a hierarchical nesting of blocks, and asserts that the Paula animation system is able to produce more natural animation with larger parent blocks compared to combinations of many




Figure 1: Score for expression (E2).

small child blocks. The system thus considers parent blocks in the hierarchy first, and falls back on developing child blocks only if no match is found.

The system accomplished this with a library of pre-animated segments and procedural techniques that correspond to blocks named after their AZee source expression. If the system recognized the expression in its entirety, the resulting animation or procedure would be invoked on the required portions of the avatar. In the absence of a match for an animation block, the system recursively falls back on animating child blocks until individual articulatory constraints would be required, sacrificing naturalness in the resulting synthesis. The multilinear nature of Paula’s skeleton integrated these animation blocks into a seamless whole even when overlapping blocks simultaneously used a given set of articulations.

3.2. Limitation for Synthesizing Proforms

What the prior work did not insist on (and indeed possibly confused by basing its narrative mostly on functional AZee trees) is that all short-cuts described and exemplified were taken based on the nesting of the score blocks, therefore only dealt with the specified *forms* to animate. The contents of blocks in the shown XML structure contain a source AZee expression which is functional, but it was used as a mere string label for a stored form with which to short-cut. This means that nothing on the functional (semantic) side of the AZee system was considered during matching. This becomes a significant issue when considering proform synthesis.

For instance, consider applying this matching scheme to the proform expression below:

(E2) `place-prf(proform-standing-person,
translate(@ST, along(medium, fwd)))`

The previous system would first run the linguistic AZee interpreter to generate an XML specification of the resulting score of nested blocks (Filhol et al., 2017), illustrated in Figure 1. It would then search for a procedural or pre-recorded animation short-cut whose name perfectly matched the source expression (E2), labelling the outer-most level of block nesting. It is unlikely that a perfect match, including the exact expression for the second argument, would be available as a pre-recorded animation due to the infinite number of possible values for this argument. The system would then fall back on the block’s child constituents (the inner blocks), which we see already consist of low-level constraints. Synthesizing from those would already be sacrificing naturalness. In this case such fallback is therefore a leap from too high a level to shot-cut, to one that is too low.

Yet the placement of a standing person does have a consistent natural dynamics, regardless of its location. There is a consistent handshape, with the upright index finger being slightly over-extended, and a consistent arcing motion, resulting from elbow and shoulder rotations which the description would specify as “downward” though it never comes out as such. If the animator has provided examples of this kind of motion, or if the system has a procedural specification for the right dynamics based on corps study, the animation system should be able exploit them, and producing in a more natural animation than one from a sparse “straight movement down” description. The “coarser the better” principle should somehow apply here despite the fact that the top-level expression cannot fully be matched. To do this, the system needs a way to look into the AZee expression and enable some form of short-cutting there. In other words, in addition to matching blocks in the *form* tree as covered by our previous bridge, we need to expand our short-cutting scheme to look into the information on the *functional* side of the input as well, i.e. the AZee expression itself.

For example, to apply a generic proform placement procedure when animating (E2), we need to recognize part of its contents, with a template like the one below:

`place-prf(proform-standing-person, X)`

where X can be anything, provided it is an expression that evaluates to a point where to place the proform, and can be retrieved for actual use with the matched animation procedure.

3.3. Proposal for new system and results

To extend the “coarser the better + fall-back” principle, the new model proposes to allow an intermediate check for matching such templates via the following extended fall-back mechanism:

1. match as label for *form* shortcut;
2. match with template for *functional* shortcut;
3. recursively process child blocks if no match (recursion terminates when reaching a block consisting of low-level constraints only).

The matching in step 2 is similar to the characterization of classifier motions as abstract templates (Liddell, 2003), with a set of parameters provided by the linguistic system. It is then up to the animation system to read and interpret those parameters. For example, consider processing the expression below, where “midssp” is a pre-defined name representing a point in the centre of the signing space.

`place-prf(proform-vehicle, midssp)`

Assuming that Paula cannot find a full match for the complete expression, the animation system will perform the following steps. Notice that we are using the word template here both for the AZee expression matching and also for the animation data specified by the artist for the specific proform *proform-vehicle*.

1. look for a template procedure for *place-prf*;
2. look for a procedure or animated version of the "proform-vehicle";
3. Evaluate the *midsp* expression to obtain the point for placement;
4. Animate the avatar using the template data.

In evaluating the *proform-vehicle* argument, the system will look for an artist specified (or motion-capture if available) template to use for the animation. Paula will then be able to leverage a range of information from that template to use in the formation of the key-frames for the action. Among these will be the handshape, the torso and clavicle parameters, the arm and elbow height as well as the articulator point on or near the arm, used for targeting the proform. In the case of the *place-prf* action above, the system would then set up two key-frames with associated velocity controls to "settle" the proform at point "midssp". This settle action is an example of one of Johnson's *ballistic* transitions that passes through the first position with a smooth speed and then eases to rest in the position of the second key. In setting up the keyframes and animating the segment, Paula will leverage the features described in section 2.2.. In particular, it will:

- read the handshape, initial elbow configuration, preferred comfortable height and other data for the proform from the artist template;
- use the IK system to set up the two keyframes for the motion (the first one with a target point that is above the final point by a small amount), adjusting the data from the artist template as needed;
- use the spine assist and livening procedures to move the avatar's torso in concert.

From these keyframes the nonlinear motion controllers will move the avatar's arms along natural arcs, which will be straight enough here to provide a perceptual "straight-down" motion.

Note that in the first item above, the system reads a name for the handshape and triggers a procedure without looking into the articulatory constraints that compose it on the child level. The second item similarly applies a generic two-keyframe layout to implement a recognized AZee pattern. We emphasize that by doing so, the system is performing a shortcut on an element of the functional expression. If at some point in this process Paula fails to find a match, for example if Paula does not have a template for the proform action, or the proform specification deviates from one of the "known" forms, the system then falls back to the expression's child blocks which will give a set of primitive constraints for the movement. Again, this would necessarily sacrifice quality but provides robustness for the system. This part of the system has not yet been integrated.

In addition to the *place-prf* template, the system currently also supports a *move-prf* template to provide movement of a proform along a path. Other proform templates will be added as the system matures. Figures 2 and 3 show two




Figure 2: A standing person moving to center




Figure 3: A vehicle moving from left to right

frames superposed from proform movements. The movement in Figure 2 corresponds to the AZee expression


```
move-prf(proform-standing-person,
         path(straight, rssp, midssp))
```

whereas the movement in Figure 3 corresponds to the LSF expression

```
move-prf(proform-vehicle, path(straight, rssp, lssp))
```



(a) place-prf(proform-standing-person, rssp)



(b) place-prf(proform-vehicle, midssp)

Figure 4: A standing person on the right of the signing space (a); a vehicle in the center (b)

The slight blur in the torso and face in 2 is indicative of the subtle spine movement that automatically supports the motion. Larger arm motions across the body as in Figure 3 are naturally accompanied by a larger torso motion, just as they would be for a live signer. This avatar torso movement is an improvement over the results of prior efforts in which the torso was stationary, and contributes to greater naturalness in the resulting animation.

Figures 4 and 5 demonstrate the generated placement movement for a selection of proforms and placement sites. The system can use any site in front of the body for placement and movement, both for one-handed and two-handed proforms. These figures further illustrate the responsiveness in the torso algorithm through the degree of blur in the head caused by the superposition of the two frames. The torso reacts more to the two-handed motion in figures 5a and 5b, causing a larger blur relative to the motion shown in figures 4a and 4b.

An animated version that demonstrates placement and movement of a selection of proforms is available online at <http://asl.cs.depaul.edu/proforms/proformPlacementAndMovement.mp4>. This animation also provides examples of placement in several locations to show the flexibility of the system. In all of these cases, Paula uses its procedural models of human motion along with data from the artist template to provide natural poses and transitions.


4. Conclusion and future work

This paper represents a first attempt at achieving more natural proform placement and movement using the structured linguistic model AZee to drive the hybrid animation system Paula. By extending the existing XML matching scheme for shortcutting to allow templating on the functional AZee expression, the system provides a flexible way for the animation system to leverage animator specified data for known proforms to improve postures and movement, while specifying a robust fallback algorithm, that still requires implementation. This sacrifices quality when the proform actions or the proforms themselves are not recognized.

Moving forward, the bridge currently supports a selection of movement procedures, but more work needs to be done both linguistically and geometrically to identify other proform structures that can be shortcut. In particular, the present work has focused on placement and straight movement of isolated one and two-handed proforms. More generally, proforms in sign languages exhibit a wide range of motion styles including bounce and wavy motions that indicate styles of movement, very general spatial movements to trace the size and shapes of objects, and relative proform placement and movement for complex scene descriptions. The current models have been built to support many of these structures, but extended study is needed for such complex motions to both refine the linguistic descriptions and naturally coordinate motion between interacting proforms.



(a) place-prf(proform-flat-round-small, rssp)



(b) place-prf(proform-flat-round-small, lssp)

Figure 5: A plate placed on the right of the signing space (a); a plate placed on the left (b)

Finally, in the future, we hope to set up a web interface providing rendered videos from input AZee expressions. This way, users will be able to connect and sandbox with our system.

Acknowledgements

The authors would like to thank Université Paris Sud for their support of Dr. John McDonald through their “recrutement de professeur invité” grant program, Dr. Annelies Braffort for her valuable input and suggestions concerning proforms and Dr. Rosalee Wolfe for her valuable input on avatar posture.

5. Bibliographical References

- Braffort, A., Benchiheub, M., and Berret, B. (2015). Aplus: A 3d corpus of french sign language. In *Proceedings of the 17th International ACM SIGACCESS Conference on Computers & Accessibility*, pages 381–382. ACM.
- 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 Braffort, A. (2006a). A sequential approach to lexical sign description. In *Proceedings of LREC 2006, second workshop on the representation and processing of sign languages: lexicographic matters and didactic scenarios*, Genova, Italy.
- Filhol, M. and Braffort, A. (2006b). Sign description: how geometry and graphing serve linguistic issues. *Theoretical issues on Sign Language research, TISLR*, 9.
- Filhol, M., Hadjadj, M. N., and Choisier, A. (2014). Non-manual features: the right to indifference. In *Representation and Processing of Sign Languages: Beyond the manual channel, Language resource and evaluation conference (LREC)*, Reykjavik, Iceland.
- Filhol, M., McDonald, J., and Wolfe, R. (2017). Synthesizing sign language by connecting linguistically structured descriptions to a multi-track animation system. In *International Conference on Universal Access in Human-Computer Interaction*, pages 27–40. Springer.
- Gibet, S., Courty, N., Duarte, K., and Naour, T. L. (2011). The signcom system for data-driven animation of interactive virtual signers: methodology and evaluation. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 1(1):6.
- Huenerfauth, M. (2004). Spatial representation of classifier predicates for machine translation into american sign language. In *Workshop on Representation and Processing of Sign Language, 4th International Conference on Language Resources and Evaluation (LREC 2004)*, pages 24–31.
- Liddell, S. K. (2003). *Grammar, gesture, and meaning in American Sign Language*. Cambridge University Press.
- Lombardo, V., Battaglino, C., Damiano, R., and Nunnari, F. (2011). An avatar-based interface for the italian sign language. In *Complex, Intelligent and Software Intensive Systems (CISIS), 2011 International Conference on*, pages 589–594. IEEE.
- López-Colino, F. and Colás, J. (2011). The synthesis of lse classifiers: From representation to evaluation. *Journal of Universal Computer Science*.
- McDonald, J., Wolfe, R., Alkoby, K., Carter, R., Davidson, M. J., Furst, J., Hinkle, D., Knoll, B., Lancaster,

- G., Smallwood, L., et al.). Achieving consistency in an fk/ik interface for a seven degree of freedom kinematic chain. In *Visual Communication and Technology Education Faculty Publications*. 3, pages 171–179.
- McDonald, J., Wolfe, R., Schnepf, J., Hochgesang, J., Jamrozik, D. G., Stumbo, M., Berke, L., Bialek, M., and Thomas, F. (2016a). An automated technique for real-time production of lifelike animations of american sign language. *Universal Access in the Information Society*, 15(4):551–566.
- McDonald, J., Wolfe, R., Wilbur, R. B., Moncrief, R., Malaia, E., Fujimoto, S., Baowidan, S., and Stec, J. (2016b). A new tool to facilitate prosodic analysis of motion capture data and a data-driven technique for the improvement of avatar motion. In *Proceedings of Language Resources and Evaluation Conference (LREC)*, pages 153–59.
- 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.
- Weast, T. (2011). American sign language tone and intonation: A phonetic analysis of eyebrow properties. *Formational units in sign languages*, pages 203–228.
- Wolfe, R., McDonald, J., and Schnepf, J. (2011). Avatar to depict sign language: Building from reusable hand animation. *International Workshop on Sign Language Translation and Avatar Technology (SLTAT)*.