Bowling Green State University

ScholarWorks@BGSU

Visual Communications and Technology
Education Faculty Publications

Visual Communications and Technology
Education

2015

Automated Technique for Real-Time Production of Lifelike Animations of American Sign Language

John McDonald

DePaul University, Chicago

Rosalee Wolfe

DePaul University, Chicago

Jerry C. Schnepf

Bowling Green State University, schnepf@bgsu.edu

Julie Hochgesang

Gallaudet University, Washington

Diana Gorman Jamrozik

Columbia College Chicago

See next page for additional authors

Follow this and additional works at: https://scholarworks.bgsu.edu/vcte_pub

Part of the [Computer Engineering Commons](#), and the [Graphic Communications Commons](#)

How does access to this work benefit you? Let us know!

Repository Citation

McDonald, John; Wolfe, Rosalee; Schnepf, Jerry C.; Hochgesang, Julie; Gorman Jamrozik, Diana; Stumbo, Marie; Berke, Larwan; Bialek, Melissa; and Thomas, Farah, "Automated Technique for Real-Time Production of Lifelike Animations of American Sign Language" (2015). *Visual Communications and Technology Education Faculty Publications*. 18.

https://scholarworks.bgsu.edu/vcte_pub/18

This Article is brought to you for free and open access by the Visual Communications and Technology Education at ScholarWorks@BGSU. It has been accepted for inclusion in Visual Communications and Technology Education Faculty Publications by an authorized administrator of ScholarWorks@BGSU.

Author(s)

John McDonald, Rosalee Wolfe, Jerry C. Schnepf, Julie Hochgesang, Diana Gorman Jamrozik, Marie Stumbo, Larwan Berke, Melissa Bialek, and Farah Thomas

An Automated Technique for Real-Time Production of Lifelike Animations of American Sign Language

John McDonald¹, Rosalee Wolfe¹, Jerry Schnepf², Julie Hochgesang³,
Diana Gorman Jamrozik⁴, Marie Stumbo¹, Larwan Berke³,
Melissa Bialek¹, Farah Thomas¹

¹DePaul University, Chicago, IL, {jmcDonald, wolfe}@cs.depaul.edu,
{mstumbo, mbialek1, fthomas}@mail.depaul.edu

²Bowling Green State University, Bowling Green, OH, schnepf@bgsu.edu

³Gallaudet University, Washington, D.C., {julie.hochgesang, larwan.berke}@gallaudet.edu

⁴Columbia College Chicago, Chicago, IL, dgorman@colum.edu

ABSTRACT: Generating sentences from a library of signs implemented through a sparse set of key frames derived from the segmental structure of a phonetic model of ASL has the advantage of flexibility and efficiency, but lacks the lifelike detail of motion capture. These difficulties are compounded when faced with real-time generation and display. This paper describes a technique for automatically adding realism without the expense of manually animating the requisite detail. The new technique layers transparently over and modifies the primary motions dictated by the segmental model, and does so with very little computational cost, enabling real-time production and display. The paper also discusses avatar optimizations that can lower the rendering overhead in real-time displays.

1. Introduction

One approach to sign language synthesis creates animations using a library of signs (Hanke, et al., 2011) (Delorme, Filhol, & Braffort, Animation Genreation Process for Sign Language Synthesis, 2009) (Efthimiou, et al., 2009) (Wolfe, McDonald, & Schnepf, An Avatar to Depict Sign Language: Building from Reusable Hand Animation, January 10-11, 2011) (Gibet, Courty, Duarte, & Le Naour, 2011). The signs are procedurally combined into phrases and sentences using tags akin to annotation tags used in linguistic analysis of sign language. In a sign language, the phonological parameters of lexical items can change depending on their usage, and new motions can be layered, depending on structural features of the discourse, such as posing questions or using constructed dialogue (Metzger, 1995) (Wilbur, 2000).

As distinct natural languages with individual grammars and structures, sign languages are used as primary languages by the Deaf worldwide. Sign synthesis is a promising method both for Deaf-hearing communication through computer translation, and also for their study and preservation. To address

these opportunities, a variety of projects have striven towards the goal of synthesizing animations of sign language.

Such efforts have exploited a range of different animation techniques from key-frame animation (Delorme, Filhol, & Braffort, Animation Generation Process for Sign Language Synthesis, 2009), to motion capture (Gibet, Courty, Duarte, & Le Naour, 2011), to procedural motion synthesis (Wolfe, Cook, McDonald, & Schnepp, 2011). A more complete review of the current literature on sign synthesis can be found in (Courty & Gibet, 2010). Many center on building sign syntheses from libraries of signs that are either transcribed by artists or captured from live sign performance. Sign language synthesis brings its own unique challenges to animation, as it often imposes rigorous requirements both on the fine detail of the motion, as in the fingers for handshapes, and on the flexibility of the motion for editing and recombination (R. Brun, Personal Communication, July 12, 2013, London, U.K.).

Both key-frame and capture techniques are expensive and time-consuming. Motion capture requires specialized skills to clean up inherently noisy data and is difficult to edit. Sign synthesis through key-frame animation requires artists who have highly trained eyes for human motion. Both techniques require experts who are also highly versant in the structure of the target sign language. Because of this, some procedural techniques are enormously appealing as they have the potential to augment and speed the work of trained animators.

The rest of this paper is organized as follows: Section 2 introduces the problem of robotic motion as it occurs in signing avatars. Section 3 describes a new automated system of spinal movement to alleviate robotic motion. Section 4 discusses optimization to aid in real-time rendering of an avatar. Section 5 covers an experiment to evaluate the spinal system and avatar optimizations, and section 6 analyzes the results. Section 7 draws conclusions and suggests future directions.

2. Challenges

One of the great challenges of character animation in general is the management of animation keys. Manipulating both the poses of an avatar and the timing of movement are greatly simplified by limiting the number of keys that describe a motion. This is especially true in sign synthesis, where modifying animation keys based on linguistic rules is easier when there are fewer keys. Such sparse animation data correlate well to the segmental structure of the Signed Language Phonetic Annotation (SLPA) parameters described in (Johnson & Liddell, 2011). Of the current representations available (Sutton, 2014) (Hanke, An overview of the HamNoSys phonetic transcription system, 2010), the SLPA model was found to be the best suited for this task.

The resulting ease of manipulation is, however, counterbalanced by a lack of realism in the animation. The very simplicity of a sparse set of keys afforded by the linguistic parameters limits the natural subtlety that one can achieve in the motion. Naturalness has long been a goal in all areas of character animation and sign synthesis is no exception. Without the subtle motions of human signing, an avatar can become highly robotic. Naturalness in motion is extremely important for the acceptability of synthesized sign.

Robotic signing, the equivalent of robotic synthesized speech, can be off-putting, can distract the viewer from the meaning, and can tire viewers in long discourse. In previous studies (Schnepp, A representation of selected nonmanual signals in American Sign Language, 2012), participants had no trouble understanding the generated sentences, however since the fine details that characterize natural motion were missing, the sentences elicited comments such as “She looks stiff” or “That’s awkward” or “I wouldn’t sign it like that”.

Ironically, linguistic tags offer little help for this problem, because their purpose is to abstract structure from the fine details of motion. The SLPA model separates sign into postural and transforming segments, analogous to the keys and interpolations of computer animation, though a naïve implementation of the model results in animation which lacks the fine details of human motion necessary for naturalness and believability in an avatar. The same abstraction that is essential for creating signed sentences as animation contributes to the awkwardness of the production.

Unfortunately, basing animation on the abstraction of the SLPA model is only one contributor to awkwardness in sentence production. The following section examines these factors in depth.

2.1 Contributing Factors to Robotic Motion

Animations procedurally created from a library of signs based on linguistic parameters naturally create a sparse distribution of keys, which makes it easy to change the animations based on linguistic context. However, that same sparseness of keys can easily lead to robotic motion. The paucity of keys can come from a variety of factors, partly due to the fact that the SLPA model is underdetermined with respect to generating animations. Other factors have to do with ignoring kinematics and animation principles. The following paragraphs examine each factor in turn.

2.1.1 Under-determination in phonetic models

Since the extra-linguistic motions are seldom specified, creating animations by relying exclusively on geometric interpretations of the designations in the SLPA model will create animations where some joints of the avatar's body remain stationary. However, when a human produces the sign, these same joints would be in motion.

For example, the ASL sign THINK only specifies the behavior of the strong hand (index finger extended and thumb and other fingers fully flexed) and contact between the index finger and the ipsilateral temple on the strong side of the body. The contact information determines the behavior of the strong arm. However, because there is no linguistically significant non-manual signal (NMS) on the spinal column, the torso and neck will not move during the production of the sign. The resulting animation is a signing “pole with arms” which does not have a natural appearance.

2.1.2 Lack of small-scale motion

In real life no part of the human body is ever truly stationary. Perlin addressed this by adding small amounts of random motion to the joints to keep the avatar alive (Perlin, 1996). However, random

motion added to animations of sign language can interfere with the subtleties in the displayed message. Moreover, the parameters of the noise must be carefully chosen or the avatar will tremble or jerk unrealistically.

2.1.3 Lack of anticipation, follow-through and secondary action

Disney animators developed a set of the core principles of animation in the 1930s that did more than create motion that roughly adhered to the laws of physics (Johnston & Thomas, 1995). These principles also communicated a character's motivation and action to an audience in a convincing manner. The principles emphasized the most important parts of the animation so that viewers could easily perceive and follow the action of the story and accept the animations as lifelike and natural. Three of these core principles of animation are particularly applicable to the generation of sign, namely anticipation, follow-through and secondary action.

In anticipation, a preparatory motion shifts the animated component slightly in the opposite direction from the intended motion before the intended motion takes place. The function of anticipation is twofold. First, it is a visual cue that alerts viewers to the upcoming motion. Since viewers are expecting to see the motion, they will more likely perceive it. Second, since the anticipatory motion draws the moving component in the opposite direction from the intended one, the overall duration and trajectory is extended, thus emphasizing and drawing more attention to the motion. Figure 1 shows three frames from an animation of the ASL sign SUCCESS. The main motion of the head is upwards. However, there is a slight anticipatory downward motion.

Fig. 1. Neutral pose, anticipatory motion, followed by intended motion in the ASL sign SUCCESS.

The animation principle of follow-through can be thought of as a complement to anticipation; it happens at the end of the movement instead of the beginning and serves a similar function to anticipation. Follow-through extends the motion, overshooting the intended final location and bouncing back. Just as anticipation prepares viewers for the motion they are about to see, follow-through visually reinforces the motion after it has occurred.

Similar to anticipation, the duration and trajectory of the entire motion is extended. Fig. 2 is a motion graph depicting wrist height in the ASL sign NOW. The initial and final postures are labeled in the graph. In the follow-through motion, the wrist position drops below the final posture and rises again to settle there. This figure is a trace of the motion of the hand in a sign transcribed by an expert animator based

on examples from multiple corpora, and was reviewed for legibility by members of the Deaf community in the Chicago area.

Fig. 2. Vertical motion in ASL sign NOW.
Initial posture is at A. B is the anticipation.
Final posture is at C. D is the follow-through.

The last animation principle, secondary action, describes movement that happens in addition to the intended motion. In ASL it occurs in two-handed signs like WORK and AGAIN where there is contact between the hands. This produces movement, however subtle, in the weak hand. Fig. 3 shows a motion trace close-up of the wrist in the weak hand during the final posture of the ASL sign AGAIN.

Fig. 3. Secondary action in weak hand in the ASL sign AGAIN.
Red boxes show weak hand/wrist motion.

These principles were first applied to computer animation in the late 1980s in computer animated short subjects such as *Luxo, Jr.* when Disney-trained artists began using computer animation packages (Lasseter, 1987). Although features such as *Cars* and *Ratatouille* are appealing and effective in conveying motion and emotion through computer-assisted hand animation, they require an enormous amount of attention to animation detail and are expensive to produce.

For sign synthesis, one must remember that these principles were generally exaggerated quite extensively for cartoon animation; however, they do describe physical processes that our bodies perform. Therefore, while not exaggerated in the way that the Disney artists originally conceived, these aspects of motion are still critical to the lifelike appearance of the movement.

2.1.4 Overuse of inverse kinematics

A technique that facilitates quick results when animating by hand is inverse kinematics (IK). IK allows the animator to specify the location of the last component, or end-effector, in a hierarchy. The rotations of each component between the first and last are automatically computed (Tolani, Goswami, & Badler, 2000). It is possible to specify various constraints in order to achieve a more natural body pose. While this technique is faster than manually adjusting each component in the hierarchy, problems arise when attempting to use IK for upper body motion. The result is often a marionette-like motion where the trajectory of the end-effector is linear, rather than a natural trajectory of an arc. Fig. 4 compares two implementations of a deictic point in ASL. The avatar on the left uses forward kinematics, not IK, for the interpolation, and the hand follows an arced trajectory. The elbow position is relatively stationary. The avatar on the right uses IK to create the motion. To maintain the linear trajectory, the elbow skates backwards in an awkward manner. The difference is marked by the yellow highlight on the elbow.

Fig. 4. FK versus IK.

In addition, IK techniques can often force joints into unrealistic configurations due to overly constraining the hand to a fixed position, e.g. vertical in front of the body. Hand positions are always somewhat approximate rather than geometrically perfect, and, especially in natural signing, comfort and efficiency will modify the orientation of the hand enough to make it comfortable but not too much so as to lose the meaning of the sign.

2.1.5 Linear Interpolation of Joint Angles

Linear interpolation causes unnatural and abrupt changes in velocity not found in nature. Linear joint interpolation leaves no room for smooth acceleration or deceleration that are the hallmark of natural human motion derived from muscular forces.

This phenomenon will appear in almost every animated sign, yielding abrupt changes in speed and direction that give a very robotic feel to the motion.

2.1.6 Synchronicity

In the transition after a hold, not all body joints begin moving at the same time (Whitaker & Halas, 2008). When turning to the side, as in a role shift, the eyes move first, followed by the head, neck, shoulders and spinal column. Further, there is a tendency for the hand to complete its transition to a new handshape well before the arm arrives at its new location as can be seen in Fig. 5 which depicts the ASL sign INFORM. Unfortunately, the SLPA model does not capture this internal timing in a sign. A naïve interpolation between postures in a sign results in a synchronous motion that is perceived as being stiff.

Fig. 5. ASL sign INFORM:
Final handshape at B; final position at C (Poor, 2008).

2.2 Alternatives to Reducing Robotic Motion

It is possible to overcome these challenges through careful manual animation, but this is a labor-intensive process requiring highly skilled artists. Moreover, the animators must have some knowledge of sign language; a rather rare combination. Results from such skilled animators are highly realistic, although time consuming and very expensive. Moreover, unless the project is structured very carefully, the linguistic components of the discourse can be buried in the kinetic motions produced by the animator as they layer on the realistic variations characteristic of human motion.

The separation of linguistic motion from kinetic motion is absolutely critical if the synthesized sign is to be able to change according to morphological rules. It also is important if the animations will be used for linguistic testing, analysis and verification. The problem is that the motion in animated signs is influenced by both linguistics and the natural kinematics of human motion. The linguistic components of a sign and its containing discourse determine the gross movements of the sign, but the kinetic motions of the body determine the naturalness of the sign's production.

This separation can be achieved by using procedural methods to aid animators not only in creating individual signs, but also in combining and modifying them to build longer discourse. The goal is to infer as much as possible by drawing on the knowledge of human kinematics to produce lifelike, convincing motion that adheres to the SLPA model of American Sign Language.

Using procedural methods shortens the development cycle by aiding animators. A new technique as part of a procedural/hybrid system is proposed, which still relies on hand animation for the sparse keys

following the segmental structure dictated of the SLPA model, though addresses several problems mentioned in the last section through procedural means. When creating an animation of an individual lexical item for the library, the linguistic components of the sign are recorded for each key to facilitate later modification and recombination. Longer discourse is then built by applying procedures dictated by linguistic tags to drive the automatic recombination of hand-animated signs. Finally, procedural techniques add the kinetic motions.

3. A New Automated System for Spinal Movement

The first part of the proposed solution deals with the lack of motion in the spine as the avatar's arms move according to the SLPA model. The proposed extension will work with any limb IK solution, either iterative or analytic (McDonald, et al., 2002), and is applied before running the IK solution itself. It assumes that the spine and collar bones of the model are not involved in the IK chain.

While iterative IK systems do allow for the extension of the IK chain through the spine to the hips of the model, these systems become increasingly slower and more difficult to control as the chain becomes longer. Other systems separate the spine, shoulder and arm into separate IK chains that can be coupled to let the spine react to the arm motion more automatically (Paolo & Ronan, 2004), (Elliott, Glauert, Kennaway, Marshall, & Safar, 2008). While they do avoid longer individual IK chains, such systems markedly increase the complexity of the computation.

The movement of the arms and torso have different purposes in the human body, and, correspondingly, in animation. The motion of the torso supports arm movement and usually precedes it. Thus, the arm motion will follow the torso temporally. Long IK chains tend to remove this timing sequence, and in fact give reverse impression that the arm is pulling on or leading the torso, resulting in the aforementioned marionette effect. Moreover, long IK chains produce a complex collection of joint rotations that cannot be easily staggered temporally to reduce synchronicity.

The application presented in this paper had three requirements that led to a more direct solution:

1. The IK system (McDonald, et al., 2002) is an analytic solution tailored to the human arm. An analytic spine solution was desired as well, so as to avoid an increased computational overhead.
2. A solution tailored to the specific type of spine motion discussed above was desired, rather than a more general spine solution. It was this fact that made requirement 1) easier to satisfy.
3. The aim was to provide a solution that allowed additional independent spine control. The solution serves as a starting point for the artist to speed workflow, but should not inhibit other scripted or automated spine motions such as role shifts.

The last of these three requirements is critical. The new extension is intended to be an assist to animators, rather than to assume sole control over the spine. Animators can still control the spine independently of this IK extension. It provides a good first estimate for the spinal movement through an automated system that will enliven the spine, eliminating one of the causes of robotic motion.

With an IK chain that solves for joint angles from the shoulder to the hand, the question becomes, “What kinds of movement will be most natural to use as a sensible default for the spine?”. Humans move their spines in a myriad of purposeful ways, and ASL is no exception to this, as spinal non-manual signals are an integral part of the language.

In the absence of purposeful spinal motion, which would normally be animated separately by the artist, one of the main functions of the spine is to extend and enhance the range of motion of the arm. Therefore, this new extension uses the reach of the arm to cue the spinal motion. The greater the distance of the reach, the more the spine will bend to assist the arm. Additionally, sequencing the animation keys involved will eliminate the perception of the arm dragging the spine in the style of a marionette. This will be discussed in detail in Section 3.5.

3.1 Movement of the Spine

When using the motion of the spine to assist the reach of the arm, the most direct solution is to use the position of the end-effector's target from the IK solution. The farther that target is from the trunk of the body the more the spine/shoulder will bend to assist the motion, and the farther the target lies across the body, the more the spine will twist to facilitate the reach. Thus, the desired position of the hand will be assisted by both a bending action that increases the range of the hand's reach, and a twisting action which facilitates the lateral movement of the hand across the body.

The magnitude of this effect is highly dependent on the distance of the arm reach, though when the hands are within the more proximal regions of sign space, there will be very little effect on the spine. This displacement should, however, increase rapidly when the hands begin to reach beyond the zone of comfort, and should taper off smoothly as one approaches the limit of the spine's rotation.

The rotation of the spine will also depend equally on both hands, so the computation must consider the positions of both the left and right IK targets symmetrically. Moreover, the interaction of the hands with the spine differs for the torso bend and the torso twist. In fact, there are at least three distinct influences which must be blended smoothly:

1. Reaching with one or both hands will bend the torso towards the target(s) and the effect will be greatest when both hands are directed towards the same target. Thus, the bend is additive for the two arms.
2. Lateral motion of a hand across the body will cause the torso to twist and rotate the shoulder towards the direction of motion. This effect is also cumulative. If the arms go in the same direction, the effect will be amplified, and if in opposite directions, they will cancel each other.
3. In addition to the torso bend, when one hand reaches out from the body, the torso will also twist to move the corresponding shoulder towards the target. This effect is different from the previous two in that equal motion between the two arms will cancel out the effect.

To build the extension, it is assumed that the coordinate system on the avatar's torso is oriented with the x-axis towards the right, the y-axis forward and the z-axis up. Consider the right and left articulator

targets A_R and A_L , and the corresponding neutral shoulder positions S_R and S_L . The displacement vectors from each shoulder are built to the corresponding target as

$$\begin{aligned} V_R &= A_R - S_R \\ V_L &= A_L - S_L \end{aligned}$$

Let V_{reach} be the average of these two vectors:

$$V_{reach} = \frac{(V_R + V_L)}{2}$$

To calculate both the bend angle and direction, the torso must be rotated in the direction of V_{reach} . The bend angle is computed by first normalizing by the arm's length, and then applying a smoothstep function (Wyvill, McPheeters, & Wyvill, 1986) to both mute its contribution near the body and smoothly clamp the effect at the end of the arm's range:

$$\theta_{bend} = c_0 * \text{smoothstep}\left(\frac{|V_{reach}|}{L_{arm}}\right)$$

where c_0 is a tuning parameter that controls the magnitude of the effect, and

$$\text{smoothstep}(x) = 3x^2 - 2x^3$$

The axis of rotation is given by the cross product of the arm direction with the vertical axis vector $Z = \langle 0, 0, 1 \rangle$.

$$a_{bend} = \frac{Z \times V_{reach}}{|Z \times V_{reach}|}$$

The matrix for a rotation of θ_{bend} about the axis a_{bend} is formed in the usual manner, and if the avatar uses Euler angles for joint control, the angles can be extracted as in (Shoemaker, 1994).

In the case of a reach up or down, there will be no effect, thus will be necessary to rely on the shoulder to extend the reach. One could extend this easily to arch the back for such a reach, however that is beyond the scope of the present work.

For the twist, a more involved combination of the two wrist displacement vectors will be needed, since the twist depends on the x and y positions of the hand differently:

$$V_{twist} = \langle V_{R,x} + V_{L,x}, V_{R,y} - V_{L,y}, 0 \rangle$$

This calculation adds in the x-direction and cancels in the y.

The twist angle is obtained by normalizing V_{twist} by twice the arm length, since the x-term is a sum rather than an average, and taking the dot product with a constant weighting vector V_w

$$V_w = \langle X_w, Y_w, 0 \rangle$$

whose components determine the amount of twist contributed by both the lateral and distal movement of the arms. This dot product then passes through the same smoothstep function as before to apply the effect gradually when the hands are in close proximity to their respective shoulders, and to smoothly clamp the effect at the far range of motion:

$$\theta_{twist} = c_1 \text{smoothstep} \left(V_w \bullet \frac{V_{twist}}{2L_{arm}} \right)$$

Here again, c_1 determines the overall size of the twist.

3.2 Spreading the Effect Over the Spine

The human spine is a chain of 24 articulating vertebrae which are seldom modeled independently in an avatar. Most often the articulation of the spine is approximated by three or more conventional rotational joints, which correspond to the major subdivisions of the spine into the cervical, thoracic and lumbar sections (Saladin, 2007). The present avatar has three joints in its back which correspond roughly to the base of the lumbar, the articulation between the lumbar and the thoracic and, finally, an articulation about 2/3 of the way up the thoracic region, creating lower, middle and upper spine joints. The cervical portion of the spine is subsumed by the neck joints which are not considered here.

For basic bending and twisting motions, the human vertebrae rotate in concert. In the model, the division of the desired rotation among the set of joints in the spine is automated by weighting each joint's rotation according to the relative range of motion in the corresponding subdivision of the spine. For these ranges, see (Reese, 2010). Thus, for example, in the spine bend, one would set:

$$\begin{aligned}\theta_{lower} &= w_0 * \theta_{bend} \\ \theta_{middle} &= w_1 * \theta_{bend} \\ \theta_{upper} &= w_2 * \theta_{bend}\end{aligned}$$

where the weighting coefficients w_i sum to 1. These coefficients will differ for the bend and twist angles since the ranges of motion in the spine are different for each of these directions. It is also important to note that for the torso bend, these weights are easiest to apply if Euler angles are used for the torso rotation since they nicely correspond to the usual ranges of motion cited in most medical texts such as flexion-extension, bending and rotation (twist).

3.3 Shoulder Movement

The movement of the human shoulder is created by two primary articulations, the acromioclavicular and sternoclavicular joints that connect the clavicle, scapula and the rib cage (Saladin, 2007). It is most often approximated in an avatar using a rotational joint placed at the intersection of the spine and the neck, and having two primary axes of rotation. The vertical rotation allows the raising and lowering of the shoulder, while the anteroposterior rotation moves the shoulder forwards and backwards. While the clavicular joints also have a small amount of rotation along the longitudinal axis to twist the shoulder, this is often ignored in favor of incorporating the motion directly in the shoulder joint itself.

The computation for automatically moving the shoulder is somewhat simpler than for the torso, due to the nature of the joint. Reaching out, in the y -direction, from the model will rotate the shoulder forward, while reaching up or down will move the shoulder accordingly. The computation follows from the same displacement vectors V_R and V_L as before, except that the y and z coordinates are considered separately for the two movements. Each is normalized and is passed through a smoothing function:

$$\theta_{R,raise} = c_2 * \text{smoothstep} \left(\frac{V_{R,x}}{L_{arm}} \right)$$

$$\theta_{R,lateral} = c_3 * \text{smoothstep} \left(\frac{V_{R,y}}{L_{arm}} \right)$$

where the c_i determine the range of the shoulder motion.

3.4 Results of the Extension

Fig. 6 displays an example of Paula, the present avatar, reaching across her body. The first image is the raw posture derived from the original IK limb computation. The second displays the more natural posture achieved by the spine and shoulder extensions described above.

Fig. 6. Reaching posture of the model. Image at left is without the IK extension. Image at right is with the extension.

3.5 Timing

The last two sections addressed one of the most important aspects of avoiding robotic behavior in the model. Maintaining motion in all of the joints in the spine alleviates the "pole-with-arms" syndrome and gives the avatar a more life-like appearance. There is, however, one aspect that is made somewhat worse by this solution when used with a set of very sparse keys. All the avatar's joints will start moving at the same time and come to rest synchronously. Asynchronous timing of linguistic components of sign has been investigated (Filhol, 2011). The work presented here extends these timing considerations to extra-linguistic motions that occur in concert with a signer's arm movements.

The human body, as a complex system of joints, bones and electrical impulses simply never moves with such precision and coordination. The joints begin moving in progression, separated by such short

intervals that the effect is not easily noticeable. However, if this progression is missing, the avatar will appear robotic. The progression used arose from a human motion study with an expert mime (Sozio, 1989), and is the following:

1. The eye-gaze shifts to a focus point
2. The neck turns to follow
3. The spine and hips shift
4. The clavicle moves
5. The arm moves

In terms of animation timing, these steps may be separated by only one or two frames, though the difference is subtly noticeable and can make a large difference in the naturalness of the motion. In fact, for IK systems, this progression is one of the main causes of the marionette effect mentioned earlier where the arm appears to drag the entire upper body. Being so used to the progression, in its absence, one tends to perceive that the arm is in fact leading the rest of the body.

The timing solution itself is a simple heuristic, but requires special handling in an animation system. Coming into or out of any posture all that is needed is to stagger the spine, shoulder and arm joints with a small amount of time between key frames. At 30fps, at least one or two frames between each key would be desired. For example, while the animator intends that the motion begin, for example, at frame 45, the system may actually start the spine motion at 43, the shoulder at 45 and the arm at 47. This can be done transparently if the keys in the animation are sparse, but is more difficult with many, tightly packed keys. The resulting effect is to break up the synchronicity of the movement and give it a more natural muscular appearance.

Going further, this progression can be broken down to a finer granularity. The motion of the spine itself is not completely synchronous either, because the hips tend to move before the upper spine as they initiate the motion. So, in a large spinal motion such as occurs in a role shift for constructed dialogue, the shift will appear much more natural if the spinal joints are staggered slightly. Since there is a much tighter relationship between the spine joints, the corresponding key distribution will also be tighter.

3.6 Handling Held Postures

An additional contribution to robotic behavior from the SLPA model is the representation of postural segments. While the model does not dictate a complete cessation of motion, it does not, and should not, dictate the autonomous motion that occurs while joints are at rest, however momentarily. Likewise, joints that are not involved in the sign, and have no linguistic specification, will remain held. An example of this is the set of joints in the weak hand and arm for one-handed signs. The problem is that the human body is simply incapable of coming completely to rest, unlike a mechanical robot. An avatar perfectly at rest is perceived as a still image and not a 3D animation.

To address this, a small component of random movement (noise) is added to each of the major joints in the avatar as in (Perlin, 1996). This causes the avatar's joints to vary slightly over time. The challenge of applying this technique is in determining the amplitude and the frequency of the noise. Noise with high

amplitude and/or frequency will give the avatar a very jittery or shaky appearance and will be worse than no noise application at all. The solution is to tune the noise so that its effect is just above the threshold of visibility, both in amplitude and frequency. The following guidelines from human anatomy are used to configure noise:

1. The amplitude of the noise should decrease for more distal joints since the muscles are generally smaller.
2. The frequency of the noise should increase for more distal joints. This can be seen, for example, in the trembling of an extended hand as one attempts to hold it still.
3. The amplitude of the noise will be greater when a part of the body is not under purposeful control. For example, in a one-handed sign, the strong hand will have random motion with much smaller amplitude than the weak hand.

4. Optimizing the Avatar for Real-Time Display

The goal of efficient real-time display has guided the formation of the techniques presented here for extending the analytic IK algorithm to the Avatar's torso. Rather than working with more traditional optimization based or iterative techniques, the present models are analytic in nature with extremely low computational overhead for the rendering system, as all the following components have computational complexity of $O(1)$:

1. The analytic IK method for the human arm
2. The extension of the IK method to the torso
3. The noise algorithm

The purpose of these techniques is to aid an animator in the time-consuming task of creating animation containing motion that is naturally lifelike. Creating such motion is a skill that requires years to master. Even for skilled practitioners, animating a sign can require many hours to complete. Automatically livening the torso of the Avatar serves as an aid to the animators, shortening the animation process. The system automatically calculates some of the base movement that the animator can then build on.

One of the overarching long-term goals of this research is to make life easier for the animator by extending these techniques and combining them with other models of human motion to automatically increase the naturalness of the generated sign. Such effort saving techniques will, however, come with an increased burden on computer hardware. Reducing the rendering complexity of the scene gives us additional processor and GPU overhead that can be used to incorporate such techniques. To this end, ways that could simultaneously lower the rendering complexity and better balance the real-time performance of the system without sacrificing the visual fidelity of the avatar have been investigated.

ASL synthesis places unique burdens on avatars due to the high fidelity required to communicate the full range of ASL's expressiveness. One of the primary challenges is the range of facial NMS which form key linguistic components in ASL at all levels. In fact, facial NMS can radically change the meaning of individual lexical items as well as whole phrases in addition to their more universal role in human communication to express affect. Because of this, the fidelity of the avatar's facial model is paramount,

and non-realtime ASL rendering systems will often ironically use more polygons for the face, mouth and tongue than they do on the arms and hands of the model. Consider the initial composition of the avatar that has been long used in this project for offline rendering. The head, hair, teeth and eyelashes contained over 70% of the avatar's polygons. In particular, the eyelashes are important for facial NMS as they emphasize the eyes.

Table 1: Polygon counts for main avatar components

Object	Polygon Count
Eyelashes	17,356
Tongue and Teeth	8,929
Hair	8,088
Head and Face	6,180
Sweater	8,300
Hands (each)	3,922

Unfortunately, such complexity presents challenges for real-time display, both in the raw number of polygons and in the number of discrete sub-objects contained in the model. Real-time systems are best optimized when large collections of polygons can be pushed to the renderer all at once (Shreiner, Sellers, Kessenich, & Licea-Kane, 2013).

The goal was thus to optimize the avatar for real-time display, though without any noticeable reduction in fidelity. The first effort, which focused on the model's hair, drives home the restrictiveness of this goal. The hair seemed a promising candidate for optimization because it is not linguistically significant and contains 12% of model's polygons. An attempt was made to optimize the polygonal meshes by merging the closest vertices using a standard mesh decimation tool. Although this optimization eliminated 6,000 polygons, the results were not judged aesthetically acceptable. The profile of the hair had a jagged appearance (Fig. 7). Since it was not possible to find an acceptable balance between polygon reduction and appearance, it was decided that the hair should be left alone turning attention to the eyelashes, the smallest objects on the model, but which used the highest percentages of the polygons.

Fig. 7. Hair before and after optimization.

4.1 Optimizing eyelashes and teeth

Eyelashes are important linguistically, because the eyes are involved in a large number of NMS, and eyelashes emphasize the perimeters of the eye apertures, making them more visible. They were initially modeled as a set of individual extruded 3D tubes. This was satisfactory in an offline rendering environment, though not practical for real-time rendering. In the current avatar, the thousands of eyelash polygons are concentrated in a very small portion of the model that in turn occupies a limited area of the final render.

To optimize the eyelashes the individually modeled lashes were replaced with four curved surfaces and used texture and opacity mapping to convey the appearance of individual lashes. The curved surface followed the contour of the original lashes (Fig. 8) and allows the lashes to be seen from all angles. A flat surface would suffer the disadvantage of billboarding where the object will disappear when seen edge-on. To create realistic variation in lash length and thickness, an artist drew an opacity map that determined which parts of the surface would be visible.

Fig. 8. Geometry and opacity map of optimized eyelashes.

The revised version of the eyelashes had only 660 polygons in comparison to over 17,000 polygons in the initial version. Additionally, the revised model replaced the somewhat regular lashes with ones that were more natural in appearance (Fig. 9).

Fig. 9. Eyelashes before and after optimization.

Similar to the eyelashes, the teeth contained a high number of polygons in a small, hard to see area with a small display area in the final render. The newly optimized version of the teeth is a curved surface mimicking the contours of the original teeth model (Fig. 10). Instead of attempting to model each tooth as in the original the appearance of teeth was created through the application of a texture map and an opacity map. Both maps were created from a composite of renders from the original model.

Fig. 10. Geometry and texture maps of optimized teeth.

Through the optimization process the number of polygons in the teeth was reduced from over 8000 to 440. For normal viewing purposes the optimized teeth appear essentially the same as the originals (Fig. 11).

Fig. 11. Teeth before and after optimization.

4.2 Lighting

Careful choices in lighting techniques can help a viewer perceive the avatar more clearly as well as speed up rendering times. Even lighting can enhance facial features and thus the legibility of facial NMS. Unfortunately, each additional light places increased computational burden on a real-time rendering system. The lighting calculations for the current avatar are based on the Phong Illumination Model (Phong, 1975) which incurs a modest computational cost that increases linearly with the number of lights. Initially, six lights were used, which provided even lighting, but slowed real-time rendering rates.

To address this, the authors had to revisit a technique in practical portrait photography called “three point lighting”, which removes unflattering shadows from the subject’s face, while emphasizing the facial features that are key to NMS. This system employs the use of a fill light and key light to illuminate the subject directly, and a third light to either provide subtle highlights or emphasize the subject’s silhouette. Usually, the key light is placed in front of the subject to serve as the primary illumination. A fill light placed to the side of the subject eliminates distracting shadows, and the third light is situated as an accent to open or close specific shadows (Napoli & Gloman, 2007).

For the current avatar, three lights were found insufficient due to the lack of simulated ambient light in the rendering system. To compensate, a fourth light was added. **Fig. 12** contains sample renderings of the original and optimized lighting. A nice side effect of the new setup was the removal of the line of symmetry previously cast along the vertical axis of the face.

Fig. 12. Comparison of lighting setups, at left initial lighting and at right optimized lighting.

A final optimization was achieved by ensuring that only the key light cast shadows, thus removing the time and memory costs of additional shadow computations.

These optimizations provided a reduction of polygons comprising the head, hair, teeth and eyelashes. These originally comprised 70% of the avatar's polygons and now comprise less than 30%. In addition, reducing the number of point-lights from 6 to 4 reduced the computational load of the lighting configuration by a third and resulted in higher fidelity lighting on the avatar (see **Fig. 12**). The optimized avatar appeared in an evaluation, described in the next section, as a check on its fidelity.

5. Evaluation

The new spine technique relieved artists from the time-consuming task of manually animating the spinal column. However, it remained to be determined if the automated spine movement enhanced either the legibility or acceptability of generated ASL. It was also interesting to know if the newly optimized avatar would be received favorably. A study was conducted to address two research questions:

- Does the automated spine system generate movement that is clear, correct, understandable, and natural?
- Does the addition of movement from an automated spine system produce animation that is preferable to animations without it?

The first question evaluates the fidelity of the polygon and lighting optimizations as well as the naturalness of the automated spinal movement system. The second question evaluates the efficacy of the underlying mathematical model for spinal motion. This study was reviewed and approved by three university Institutional Review Boards (DePaul IRB #RW071813CDM, Columbia College IRB #2014-00167, Bowling Green State University IRB #507016-2).

5.1 Stimuli

The test stimuli were sentences from a complete story about a childhood memory as listed in Fig. 13. Fig. 14 lists the five sentences in the stimuli as glossed ASL. A complete story is necessary to provide context for such linguistic processes as indexing, agreement verbs, and role shifts. Without the proper context, these lose their meaning. All of the sentences were grammatically correct. All verbs were conjugated and all NMS were present, including syntactic markers. In addition, the sentences contained appropriate pragmatics.

A panel of language experts reviewed several versions of the sentences before approving their final form. The panel consisted of three nationally certified interpreters and two native ASL users, including a linguist.

I remember when I was small, I lost my little toy car. I looked everywhere for it. I was so sad. But, surprise! My brother bought me a new one. It was big and blue. Blue is my favorite color. I was so happy!

Fig. 13. Story for Evaluation Study (English).

1. LOOKING-BACK REMEMBER ME SMALL, HAVE LITTLE T-O-Y C-A-R.
2. I LOST WHERE. I SEARCH, SEARCH, CAN'T FIND -- SAD.
3. BUT SURPRISE -- MY BROTHER GET NEW C-A-R, BIG, BLUE, GIVE-ME.
4. INFORM-YOU BLUE MY FAVORITE COLOR.
5. FINALLY HAVE C-A-R AGAIN -- HAPPY!

Fig. 14. Story for Evaluation Study (glossed).

Two versions of each sentence in the story were generated. The control version did not have the automatically-generated spinal movement and the treatment the second version did. All of the manually-animated poses were present in both versions. If the artist had included an overt spinal pose, it was included in both versions of the animation. In addition, both versions included noise, so the spine was not entirely stationary, even in the control stimuli. Only the automated aspects of spinal movement were missing from the control stimuli. The stimuli are available for viewing at <http://asl.cs.depaul.edu/video/McDonaldUAIS2015.avi>.

5.2 Participants

Participants were all adults fluent in ASL. The invitation to participate was posted on several Deaf newsgroups as a link to a video, which contained the fingerspelled URL address for the test.

5.3 Procedure

The tests were performed online, using SignQUOTE (Schnepp, Wolfe, Shiver, McDonald, & Toro, 2011), which facilitates online testing exclusively in ASL. Participants used their own computer, equipped with a web cam, to complete the test.

Participants viewed the informed consent, and indicated their consent by continuing on with the study and filling out a brief pre-test questionnaire. In the first section of the test, they viewed the story, one sentence at a time, and rated each sentence. They could view a sentence as many times as they wished. All of the sentences in this portion of the test included the automated spine movement (See Fig. 15 for a screen shot of the SignQUOTE interface).

Fig. 15. Test stimulus for part 1.

In the second section of the test, participants again viewed the story, one sentence at a time, but saw two versions of the sentence, presented side-by-side, as in Fig. 16. To avoid experimental bias, the positions of the two stimuli were varied, so the control sometimes appeared on the left and sometimes on the right. Participants could view each of the two stimuli as many times as they wished, as both stimuli had their own play/pause/rewind buttons.

Fig. 16. Stimulus for part 2.

5.4 Measures

For the first section of the test, participants rated each sentence for clarity, grammaticality, understandability and naturalness on a five-point Likert scale. **Error! Reference source not found.** Table 2 contains a back translation of the questions and the rating scales. For the second section of the test, participants clicked the button corresponding to the animation they preferred.

Table 2: Back translation of questions and rating scales for part 1

1. Rate the clarity of the sentence.	
Very confusing	☐ ☐ ☐ ☐ ☐
Very clear	
2. How grammatical was the sentence?	
Full of errors, unacceptable	☐ ☐ ☐ ☐ ☐
Perfectly correct	
3. How understandable was the sentence?	
Very hard	☐ ☐ ☐ ☐ ☐
Very easy	
4. How natural was the motion?	
Robotic	☐ ☐ ☐ ☐ ☐
Very human-like	

5.5 Results

In total, 22 participants completed the test. The demographic breakdown was as follows:

- 19 deaf, one hard-of-hearing, and two hearing participants

- All participants identified either ASL (16 participants) or Pidgin Signed English (6 participants) as their preferred language
- 17 had used ASL all their lives, 3 had used ASL between 5 and 15 years, and two between 1 to 2 years
- 17 of the 22 were born deaf

For part 1 of the test, participants viewed each of the five sentences in the story in turn and answered the questions in **Error! Reference source not found.** Table 2. The results for the four questions are displayed in Fig. 17 as Tukey box-and-whisker plots. The median value in each sentence is displayed as a dark-blue circle, the range between the first (25%) and third (75%) quartiles is displayed as a thick box, and the overall range is given by the whisker lines. The data in this table show that:

- With the exception of sentence 3, over 75% of participants rated clarity as a 3 or above, with 50% rating them as clear to very clear
- With the exception of sentence 3, over 75% of participants rated grammaticality as a 3 or above, with 50% rating them as correct or perfectly correct
- With the exception of sentence 3, over 75% of participants rated understandability as a 3 or above, with 50% rating them as easy or very easy to understand
- Over 50% of participants rated every sentence as human-like or very human-like.

Fig. 17. Results from part 1 of test.

Error! Reference source not found. Table 3 displays the responses from part 2 of the test as A/B preferences. In the test instrument, the A/B order of the control stimulus varied, however the responses given below have been re-ordered to make the data more readable. The last row of the table contains the results of a binomial test expressing the probability that the results could be due to random chance. In all cases, participants preferred the spine movement at least as well as the control, and in three of the sentences (2, 3, 4) their preference for the spine movement was statistically significant ($p < 0.05$).

Table 3: Preferences from Part 2 of test

	S1	S2	S3	S4	S5
Control	11	4	6	3	9
Spine Movement	11	18	16	19	13
Binomial p-value	0.5840	0.0021	0.0262	0.0004	0.2617

6. Discussion and Conclusions

The data from the first part of the test address the research question "Does the automated spine system generate movement that is clear, correct, understandable, and natural?". Participants perceived the majority of the sentences in a positive light. The results were particularly strong for clarity and understandability where over half of the participants rated all sentences as either clear or very clear and easy or very easy to understand¹. Also, sentences received favorable ratings in grammaticality with the exception of sentence 3. This may be due in part to the fact that the structure of this sentence was more complex. Sentence 3 was the only one where the signer was not the subject of the sentence, and the only one that contained both direct and indirect objects. These results compare favorably to results from similar user studies of generated ASL (Kacorri, Lu, & Huenerfauth, 2013) (Huenerfauth, Lu, & Rosenberg, 2011).

English: My brother bought me a new one. It was big and blue. ASL: BUT SURPRISE -- MY BROTHER GET NEW C-A-R, BIG, BLUE, GIVE-ME.

Fig. 18. Sentence 3 of the test stimuli.

Overall, the ratings for naturalness were lower than for the other three measures with 50 percent of the respondents rating all of the sentences as neutral or above. It may be that naturalness is the most demanding criterion of all. This is analogous to synthesized speech. It is possible to *understand* a *grammatical* synthetic voice message *clearly*, such as an automated airline flight notification sent to a person's phone, but it would not be perceived as *natural*. Although lower than the other three, this result still compares favorably with prior tests of naturalness in key-frame and procedurally generated animations of sign language avatars.

Data from the second part of the test yield insights into the second research question, "Does the addition of movement from an automated spine system produce animation that is preferable to

¹ Recall that a median of 3.5 on an integer 5-point Likert scale indicates that 50% were either a 4 or a 5.

animations without it?". In the majority of cases, the answer appears to be "yes". Responses from participants showed a statistically significant preference ($p < 0.05$) for the added spinal movement sentences 2, 3 and 4.

For sentences 1 and 5, however, the responses showed no statistically significant preference for either version, although a majority of participants preferred the added spinal movement in sentence 5. This may be due to the fact that, in these sentences, the spinal motion caused by the system was largely in the forward-back direction. In contrast, the other three sentences had more lateral and twisting motion, which is more apparent from a front camera view. Another contributing factor in sentence 1 may have been its placement as the first video in the sequence. Participants may have needed acclimation to the side-by-side video comparison. Further study is indicated, and perhaps in future tests a warm-up exercise may be appropriate.

It is notable that, in all the cases that involved lateral motion of the body, a significant majority of participants judged the automated spine algorithm favorably. Recall that the automated spine system is not intended to be a sole solution for spine animation, but an aid to animators, or a supplement to other procedurally generated techniques. This user test indicates that the system would be highly beneficial in both of these roles.

7. Summary and Future Work

The techniques described in the present work address two of the many causes of robotic motion that can arise from implementing the SLPA model of ASL in procedurally-based synthesis systems.

First, this new approach can serve as an extension to any limb IK system, and aids artists by automatically rotating the torso and spine of an avatar to support the specified arm motions from the linguistic model. The approach staggers the timing of the keys to reduce the marionette effect arising from completely synchronous joint movement. The second technique addresses the complete lack of motion in held joints. Careful applications of Perlin noise to rotational joints can enliven the avatar even when the linguistic model specifies that an avatar hold a pose, or a joint's motion is unspecified. These techniques have the potential for both aiding manual animation and also for supplementing procedurally generated avatar movements systems such as (Hanke, et al., 2011), (Delorme, Filhol, & Braffort, Animation Genreation Process for Sign Language Synthesis, 2009).

Both of these techniques still present opportunities for further study. For example, it is possible that basing the shoulder and torso extension on the position of the avatar's elbow rather than the IK end-effector would yield more realistic motions. The elbow's position is more closely related to the shoulder, having only a single bone segment between them.

Moreover, both the torso and the noise methods presented in this paper would benefit greatly from a detailed study of motion-capture data and application of a comfort model similar to the one presented in (Delorme, Modelisation du squelette pour la generation realiste de postures de la lange de signes francaise, 2011). Such a model has the potential to increase the naturalness of the generated spine motions by incorporating data from captured human motion. The present model was formed through

close collaboration between computer scientists and animators, though could benefit from both improvement and validation through a data-driven study.

The authors are studying the other causes of robotic motion mentioned here, and intend to address them in the future.

Acknowledgements. The authors would like to thank the Deaf experts and qualified interpreters for reviewing drafts of the animations. We are also grateful to the reviewers for their thoughtful comments and valuable feedback.

References

- Courty, N., & Gibet, S. (2010). Why Is the Creation of a Virtual Signer Challenging Computer. In R. Boulic, Y. Chrysanthou, & T. Komura (Ed.). (pp. 290-300). Utrecht, The Netherlands: Springer.
- Delorme, M. (2011). Modélisation du squelette pour la génération réaliste de postures de la langue de signes française. Ph.D. Dissertation, Laboratoire d'Informatique pour la Mécanique et les Sciences de l'Ingénieur (LIMSI).
- Delorme, M., Filhol, M., & Braffort, A. (2009). Animation Generation Process for Sign Language Synthesis. (pp. 386-390). *Advances in Computer-Human Interactions, ACHI '09*.
- Efthimiou, E., Fotinea, S.-E., F., Vogler, C., Hanke, T., Glauert, J., Bowden, R., Segouat, J. (2009). Sign Language Recognition, Generation, and Modelling: A Research Effort with Applications in Deaf Communication. *Universal Access in Human-Computer Interaction, Addressing Diversity. Lecture Notes in Computer Science., Volume 5614*, pp 21-30.
- Elliott, R., Glauert, J., Kennaway, J., Marshall, I., & 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. (2011). A combination of two synchronisation methods to formalise sign language animation. *Proceedings of the 9th international Gesture Workshop*.
- Gibet, S., Courty, N., Duarte, K., & Le Naour, T. (2011). The SignCom system for data-driven animation of interactive virtual signers: Methodology and evaluation. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 1(1), 1-26.
- Hanke, T. (2010, June 14-16). An overview of the HamNoSys phonetic transcription system. Retrieved December 28, 2011, from Sign Linguistics Corpora Website: http://www.ru.nl/publish/pages/570576/slc3_2010_hanke.pdf
- Hanke, T., Matthes, S., Regen, A., Storz, J., Satu, W., Elliott, R., Kennaway, R. (2011). Using Timing Information To Improve the Performance of Avatars. Dundee, Scotland, UK: Second International Workshop on Sign Language Translation and Avatar Technology (SLTAT).
- Huenerfauth, M., Lu, P., & Rosenberg, A. (2011). Evaluating importance of facial expression in American Sign Language and pidgin signed English animations. *The proceedings of the 13th international ACM SIGACCESS conference on Computers and accessibility* (pp. 99-106). New York: Association For Computing Machinery.
- Johnson, R. E., & Liddell, S. K. (2011). A Segmental Framework for Representing Signs Phonetically. *Sign Language Studies*, 11(3), 408-463.
- Johnston, O., & Thomas, F. (1995). *The Illusion of Life: Disney Animation*. New York: Random House (Disney Press).
- Kacorri, H., Lu, P., & Huenerfauth, M. (2013, October). Effect of Displaying Human Videos During an Evaluation Study of American Sign Language Animation. *ACM Transactions on Accessible Computing (TACCESS)*, 5(2), 4.
- Lasseter, J. (1987). Principles of traditional animation applied to 3D computer animation. *SIGGRAPH '87 Proceedings of the 14th annual conference on Computer graphics and interactive techniques* (pp. 35-44). Anaheim, California: Association of Computing Machinery.
- McDonald, J., Alkoby, K., Carter, R., Christopher, J., Davidson, M., Ethridge, D., Wolfe, R. (2002). A Direct Method for Positioning the Arms of a Human Model. *Proceedings of Graphics Interface*, (pp. 99-106). Calgary, Alberta Canada.
- Metzger, M. (1995). Constructed Dialogue and Constructed Action in American Sign Language. In C. Lucas, *The Sociolinguistics in Deaf Communities* (pp. 255-271). Washington, D.C.: Gallaudet University Press.

- Napoli, R., & Gloman, C. (2007). *Scene Design and Lighting Techniques: A Basic Guide for Theatre*. Burlington, MA: Focal Press.
- Paolo, B., & Ronan, B. (2004). An Inverse Kinematics Architecture Enforcing an Arbitrary Number of Strict Priority Levels. *The Visual Computer*, 20(6), 402-417.
- Perlin, K. (1996). A system for scripting interactive actors in virtual worlds. *Proceedings of ACM SIGGRAPH 96* (pp. 205-216). New Orleans: Association for Computing Machinery.
- Phong, T. (1975, June). Illumination for computer generated pictures. *Communications of the ACM*, 18(6), 311-317.
- Poor, G. (2008). ASL Video Dictionary and Inflection Guide. Retrieved May 15, 2011, from <http://www.ntid.rit.edu/dig>
- Reese, N. B. (2010). *Joint Range of Motion and Muscle Length Testing*. St. Louis, MO: Elsevier.
- Saladin, K. (2007). *Human Anatomy*. New York, NY: McGraw-Hill.
- Schnepp, J. (2012). *A representation of selected nonmanual signals in American Sign Language*. DePaul University: ProQuest Dissertations and Theses.
- Schnepp, J., Wolfe, R., Shiver, B., McDonald, J., & Toro, J. (2011). SignQUOTE: A Remote Testing Facility for Eliciting Signed Qualitative Feedback. *Second International Workshop on Sign Language Translation and Avatar Technology (SLTAT)*. Retrieved April 23, 2012, from <http://vhg.cmp.uea.ac.uk/demo/SLTAT2011Dundee/4.pdf>
- Shoemake, K. (1994). Euler angle conversion. In P. Heckbert, *Graphics Gems IV* (pp. 222-229). San Diego, CA: Academic Press.
- Shreiner, D., Sellers, G., Kessenich, J., & Licea-Kane, B. (2013). *OpenGL Programming Guide: The Official Guide to Learning OpenGL, Version 4.3 (8th Edition)*. Upper Saddle River, NJ: Addison-Wesley Professional.
- Sozio, S. (1989). *The Mastery of Mimodrame: An In-Depth Study of Mime Techniques*. Shippensburg, PA: Destiny Image Publishers.
- Sutton, V. (Ed.). (2014). *SignWriting for Sign Languages*. Retrieved October 8, 2014, from [SignWriting.org: http://www.signwriting.org/](http://www.signwriting.org/)
- Tolani, D., Goswami, A., & Badler, N. (2000). Real-Time Inverse Kinematics Techniques for Anthropomorphic Limbs. *Graphical Models*, 62, 353-388. doi:10.1006/gmod.2000.0528
- Whitaker, H., & Halas, J. (2008). *Timing for Animation*. Burlington, Massachusetts: Focal Press.
- Wilbur, R. (2000). Phonological and Prosodic Layering of Nonmanuals in American Sign Language. In K. Emmorey, H. L. Lane, U. Bellugi, & E. Klima (Eds.), *The Signs of Language Revisited: Festschrift for Ursula Bellugi and Edward Klima* (pp. 213-241).
- Wolfe, R., Cook, P., McDonald, J., & Schnepp, J. (2011). Linguistics as structure in computer animation: Toward a more effective synthesis of brow motion in American Sign Language. *Nonmanuals in Sign Language Special issue of Sign Language & Linguistics*, 14(1), 179-199.
- Wolfe, R., McDonald, J., & Schnepp, J. (January 10-11, 2011). An Avatar to Depict Sign Language: Building from Reusable Hand Animation. Berlin, Germany,: International Workshop on Sign Language Translation and Avatar Technology (SLTAT).
- Wyvill, B., McPheeters, C., & Wyvill, G. (1986). Animating soft objects. *The Visual Computer*, 2, 235-242.