

Physical Segments and Functional Gestures

Irene Appelbaum
University of Montana

1. Introduction

Co-articulation is a pervasive and widely accepted phenomenon in speech production. While phonetic segments are discrete and serially ordered, the articulators implementing them in spoken language notoriously do not operate in a corresponding fashion. In a phonetic representation of the word 'mat', for example, the segments [m], [a], and [t] are linearly concatenated and the boundaries of each segment are unambiguously identified: point to any portion of the representation and we can easily tell whether it is part of the [m], part of the [a], or part of the [t]. No portion of any one segment is part of any other. By contrast, the articulators used to produce a token utterance of the word 'mat' operate simultaneously. For example, the velum is opened and the lips closed simultaneously in the production of [m]. Obviously there is no one-to-one correspondence between individual segment and individual articulator. But nor is there a one-to-one correspondence between an individual segment and an individual *set* of articulators. Any set we try to isolate at a given time may simultaneously be implementing more than one segment. It is in this sense that speech production exhibits co-articulation.

Familiar as the above facts about co-articulation are, it is worthwhile emphasizing that co-articulation is not a claim about the operation of the articulators *per se*, but about their operation in *relation* to a particular choice of linguistic primitive -- the segment. Articulators are in motion and overlap temporally, but it is phonemes that are co-articulated.

The concept of co-articulation thus highlights a more general feature of the relation between units at the phonological level and units at the physical level: a physical description of speech -- whether in articulatory or acoustic terms -- invokes dynamic, continuous, and overlapping properties which are not isomorphic to the discrete, serially-ordered segments at the linguistic (i.e., phonological) level.

This mismatch between primitives at the phonological and physical levels has long been a thorn in the side of speech perception theorists because of the widely accepted, if often implicit, assumption that an adequate explanation requires identifying a one-to-one correspondence between units of phonological structure and units of physical structure. Logically, there are two ways to remedy this mismatch: (i) by trying to identify discrete physical segments corresponding to linguistic segments, or (ii) by trying to identify dynamic linguistic properties corresponding to the dynamic physical properties.

The most common strategy for trying to address this problem has been to adopt (i): try to identify some physical unit that *is* isomorphic to phonetic segments. That is, if the relation between individual segments and isolated articulators is not transparent, the assumption has been that there must be a set of other physical or physiological units, which is transparently related to phonetic segments. This perspective is well captured in the following quote from Liberman & Mattingly: "...invariant gestures of some description there must be, for they are required, not merely for our particular theory of speech perception, but for *any* adequate theory..." (1985, p 3). Examples of such candidate invariants have included neural structures of the motor theory (also called 'gestures') and Blumstein and Stevens' acoustic invariants.

Nevertheless, this strategy has proved wanting -- in terms of identifying even candidate invariants, in terms of empirically confirming them, in terms of identifying them for classes of segments beyond stop-consonants, and in terms of showing that these features are *used* in phonemic perception.

2. The Approach of Articulatory Phonology

More recently, the theory of articulatory phonology has approached the problem from the other direction. Instead of beginning with the assumption that segments are the phonological primitives, and searching for lower-level physical correlates, articulatory phonology begins with a physical description of speech -- specifically, an articulatory description -- and then posits higher-level descriptions of dynamic articulatory properties as phonological units.

These phonological units are dubbed *gestures* and it is a central claim of articulatory phonology that such gestures function as the primitives of phonology. As they plainly state:

In articulatory phonology, the basic phonological unit is the *articulatory gesture* (1995, p. 181).

In articulatory phonology, the basic units of phonological contrast are called *gestures* (1992, p. 155).

The basic units [of phonology] are dynamically defined articulatory gestures (1990a, p. 300).

It is important to distinguish the gestures of articulatory phonology from the traditional usage of the term 'gesture'. Gestures of both sorts are used to characterize the formation of local constrictions within the vocal tract, but there are crucial differences. First, in traditional usage gestures characterize the movement of articulators at a gross level: that is, one gesture might consist of lip closure, another of the tongue-tip against the alveolar ridge; and lip, tongue-tip, and alveolar ridge are all articulators (though a passive one in the case of the alveolar ridge).

The gestures of articulatory phonology do not characterize the movement of individual articulators, but rather the movement of "tract variables" Examples of tract variables include: LIP APERTURE, TONGUE TIP CONSTRICT LOCATION, TONGUE TIP CONSTRICT DEGREE, TONGUE BLADE CONSTRICT LOCATION, TONGUE BLADE CONSTRICT DEGREE. Although tract variables are also described using terms for gross articulators -- lip, tongue-top, etc. -- what is being specified is the degree of constriction and movement of these articulators. Moreover, unlike traditional articulatory descriptions in which degree of constriction is specified in terms of a set of discrete choices: stop, fricative, etc., the gestures of articulatory phonology are specified in continuous fashion. Specifically, tract-variable regimes specify "values for the dynamic parameters of stiffness, equilibrium position and damping ratio" (1995, p. 183).

Significant as the above differences are, the most significant departure from traditional usage is that the gestures of articulatory phonology are units of functional, that is, phonological structure. Indeed, the above description of articulatory gestures in terms of tract variable regimes is really a description of their role as linguistic/phonological elements. What the tract variable regimes specify in physical terms are also gestures -- but at a much more detailed lower level of description.

Nevertheless proponents of articulatory phonology do not present their view as containing two sets of entities -- a set of phonological gestures and a set of physical gestures. Rather, there is just a single set of physical gestures which constitute the speech event and which can be described in microscopic terms or in more abstract macroscopic terms. The distinguishing feature of the theory of articulatory phonology is the claim that one of these descriptions -- the macroscopic description -- constitutes the linguistic/phonological level description as well. The primitives of the physical description -- individual gestures -- are the primitives units of the phonological-level description as well.

Thus, in traditional articulatory phonetics gestures correspond to the physical implementation level only and are implementations of quite different sorts of units, namely, segments. In articulatory phonology gestures are descriptions of the physical implementation level, but *what* is being implemented -- the linguistic primitives -- are themselves gestures, rather than segments.

Another important feature of phonological gestures is that although they characterize dynamic properties -- indeed are specified in terms of dynamic systems -- phonological gestures are themselves time-free, discrete, and invariant. Again, Browman and Goldstein are explicit on this point:

The gestures are defined in a context-independent fashion (1995, pp. 181-182).

[Gestures are] invariantly specified phonological units (1992, p. 165).

[G]estures are defined in terms of the underlying *dynamic* that serve to characterize the motions. Such a dynamical description provides a representation that is itself time-free... (1990a, p. 300).

That is, although the physical implementation of a particular gesture may vary quite a lot depending on context, the macro-description of this gesture -- in terms of tract-variable regimes -- is said to be invariant. As phonological primitives, gestures are inherently dynamic and hence naturally match up with articulatory gestures -- the physical primitives. But like segments, gestures are discrete. Hence they also seem well suited to the role of phonological primitives -- namely, being able to account for the sounds and the sound contrasts of the world's languages by combining a small number of discrete primitives.

Although phonological gestures are discrete and invariant they overlap with each other. A word or an utterance is modeled not as a linear sequence of gestures but as a set of overlapping gestures corresponding to the overlap at the physical implementation level. Such a set of overlapping gestures, Browman and Goldstein refer to as an "ensemble" or "constellation":

[A]n utterance is modeled as an ensemble, or constellation, of a small number of potentially overlapping gestural units (1995, p. 181).

In this system, utterances are organized ensembles (or *constellations*) of units of articulatory action called *gestures* (1995, p. 182).

Articulatory phonology thus seems to be an elegant solution to long-standing problem of the mismatch between speech understood as a series of discrete functional primitives and speech understood as a continuous, dynamic physical event. By decomposing speech production in terms of its motor actions articulatory phonology accounts for both functional/phonological primitives and articulatory/physical primitives with one set of parameters, thus eclipsing the problem of how the primitives at these two levels are related. On this view, co-articulation describes, not in the first instance the relation between phonological primitives and their articulation, but in effect, the relation among individual phonological primitives, which are now viewed as themselves, simultaneous and overlapping -- like their lower-dimension counterparts, and unlike phonological segments.

3. Critique of Gestural Primitives Claim

Upon closer examination, however, the theory of articulatory phonology proves wanting in crucial respects. Two central claims of the theory are that articulatory gestures are the phonological primitives, and that the main advantage of such a choice for primitives is that it unifies the physical and phonological levels of description. In this section, I shall argue that the first claim is not adequately supported by the theory. First, Browman and Goldstein argue that gestures can be used to signal lexical differences, but they do not argue that they are the best units for the job and hence they do not provide an argument for treating them as the basic or primitive phonological units. In addition, they do not explain the relative merits of gestures over segments in terms of this functional role, and they do not show how gestures and segments can be part of a single particulate structure. Finally, even if we ignore these first two criticisms, the theory still faces a dilemma insofar as phonological gestures are said to be "defined" as invariant, whether they are so in fact or not.

3.1 Phonological Units vs. Phonological Primitives

A major task of the primitive units in a phonological theory is to signal lexical contrast. More specifically, the primitives ought to be the smallest units capable of systematically signaling meaning differences. It is not enough to show that two lexical items differ in terms of some unit. Two words

may for example differ in only one syllable, but this doesn't show that syllables are the primary phonological units. For two words that differ in one syllable will also differ in at least one phoneme. And phonemes are smaller than syllables. But in a system of syllables with no repeated phonemes -- e.g., with no minimal pairs -- the syllable, though larger than the segment, would be a better choice because it captures the smallest *systematic* difference between two lexical items. The attempt to capture lexical differences in terms of phonemes would always leave out some difference (i.e. the other phoneme in the case of a CV syllable).

Browman and Goldstein identify gestures as units that are capable of signaling lexical contrast; moreover, they repeatedly identify gestures as "basic" or "primitive" units of such contrast. Nevertheless, they do not focus their attention on justifying gestures on functional grounds -- i.e., as the units best suited to capturing such contrast. The reasoning they provide for this choice is simply that gestures are sufficient to distinguish two lexical items:

Gestures can function as primitives of phonological contrast. That is, two lexical items will contrast if they differ in gestural composition (1992, p. 57).

But as we have just seen, this is too weak a standard. Showing that two lexical items will differ if they differ in unit *x*, does not show that unit *x* functions as the basic or primitive unit in the theory.

It is interesting that the one place Browman and Goldstein *are* attentive to the difference between identifying *a* unit of phonological contrast and identifying the best one, is in response to Pierrehumbert. What is interesting is that they apply this standard not in order to justify their own choice of gestures as phonological primitives, but to criticize the choice of segments for this role:

From this perspective, then, phonemes become one particular ... solution to the necessity of capturing ... distinctive aspects of lexical items... It is, however, not a necessary solution: these same facts can be captured with other units, including gestures... (1990b, pp.418-419).

Indeed. Of course this charge could be leveled against their own gestural solution as well. But they go on to actively criticize segments as the choice of phonological primitives. Here they rightly emphasize that the choice of phonological primitive must be based on functional grounds. As they put it, the basic units should be chosen because of "their correspondence to important informational units of the phonological system" and because they are "relevant to representing theoretically significant aspects of the [phonological] system" (1990b, p. 418).

In their view, segments fare badly on this score. To explain the fact that the segment has nevertheless been a popular choice for the basic unit of phonology, they point to a confusion between practical and theoretical goals. A segmental description -- in the form of a phonetic transcription -- has been an invaluable tool for recording the sounds of the world's languages in a form that can be reconstructed by a listener. In this sense, phonetic transcriptions might be thought of as a type of 'storage device'. This does not imply, however, that they are anything more than a convenience and, in particular, it does not imply that a form of representation that works for information storage and retrieval has any real, theoretically justified status.

There is I think a significant problem with Browman and Goldstein's argument here. They distinguish the justification of the use of segments for practical goals, such as transcription, from justifying them in terms of the goals of phonological theory. It is fine to make this distinction generally, but in the case of phonology as we have seen and as Browman & Goldstein acknowledge, a central goal is to be able to systematically signal differences among lexical items. And here we know segments function very well for this purpose -- even if there is debate about whether they are the best unit, there is overwhelming evidence that they are a very good unit for the central goal of signaling lexical contrast. (Although many theories in phonology have proclaimed the death of the phoneme, segments are still used.)

So by specifying the importance of justifying phonological primitives on theoretical grounds, without specifying that such grounds include centrally being able to signal lexical contrast efficiently, Browman and Goldstein, dismiss segments without even addressing their suitability to play this theoretical role, in particular without addressing the relative suitability of segments and gestures. In the

next section, I will compare the relative merits of segments and gestures for playing a related but more general role that any primitives of the sound structure of language must meet.

3.2 *Gestures and the Particulate Principle*

To say that a phonological theory is a functional theory is to say that it tells us, among other things, how we use the sounds of language to create lexicons. Language exhibits particulate structure: enormous diversity is achieved from combining and permuting a limited number of discrete primitives. English, for example, contains about forty phonemes. If each phoneme constituted one word, the vocabulary of English would be limited to forty words; more generally, in such a case a language could have only as many words as it had phonemes. Alternatively, if words were formed by "blending" phonemes, a language would contain even fewer words than the number of phonemes. For as the number of phonemes used to form words increased, differences among words would decrease. But, as it is, words are formed by combining and permuting discrete phonemes without destroying the identity of individual phonemes – that is, words are formed by observing the particulate principle. In such a case, the lexical possibilities grow exponentially: from the forty phonemes of English hundreds of thousands of words are generated.

The particulate principle seems to be responsible for the unbounded diversity observed in biological and chemical structure as well. The millions of proteins human beings create and make use of, for example, are all composed of different combinations of only two dozen amino acids. (These two dozen amino acids, in turn, are all coded for by combining just four nucleotide bases into different three-base sequences.) Analogously, combinations of the 108 atomic elements in the universe yield an open-ended variety of molecules. Indeed, particulate structure seems to be the only (mathematically possible) way to achieve such diversity.

An additional feature of particulate structure is that it is hierarchical or compositional. That is, units at successively higher levels stand in a many-to-one relationship to each other. Individual units at a higher level are composed of a discrete number of lower-level units. In the case of biological structure each of some hundred thousand genes is composed of some combination of 20 (triplet) codons. These 20 (triplet) codons in turn are each composed of sequences of just 4 nucleotide bases.

When we turn to the sound structure of language, the particulate principle is well served by treating traditional phonemic segments as primitives. In English, for example, each of hundreds of thousands of words is composed of some combination of just 40ish discrete phonemes. If instead we treat the gestures of articulatory phonology as the primitives of the sound structure of language, we can say that all the words of English (and every other language) are composed of combinations of a small set of gestures (wide pharyngeal, closed alveolar, wide glottis, closed labial, wide velum, etc.)

Thus far, the gestures of articulatory phonology seem to be adequate candidates for the primitives of a phonological theory – at least there seems to be no *structural* problem with the claim. But if we now ask about the relation between gestures and segments in articulatory phonology, things get messier. There seem to be two main possibilities for describing this relationship: (i) either segments are compositionally derived from gestures, or (ii) segments do not play a systematic role in the description of the sound structure of language. In the latter case, much more explicit argumentation is necessary that gestures are better, more useful for describing the sound structure of language. In the former case, which is the view implied in Studdert-Kennedy and Goldstein (2003), it's problematic to claim that gestures and phonemes are hierarchically related. It may be the case that words are composed of discrete sets of gestures as required by the particulate principle, but phonemes surely are not.

So more needs to be said about how articulatory phonology can accommodate gestures as phonological primitives in a (particulate) theory that includes phonemes.

3.3 *Definitional vs. Empirical Invariance*

Assuming now that there is no problem justifying gestures as the appropriate units for phonological primitives, there is still a question about whether gestural primitives can actually do the work they are meant to, in terms of solving the lack of invariance problem. At first blush, this would

seem to be a straightforwardly unproblematic aspect of articulatory phonology, for as we have seen phonological gestures are descriptions of physical gestures at a macroscopic level. Indeed a central motivation for positing gestures as phonological primitives is just that they can't help but be isomorphic to gestures understood as units of articulatory action.

Still there is a problem that needs to be clarified. The physical implementations of the gestures can vary a great deal, but the phonological gestures are specified as dynamic systems and are held to be invariant:

The physical properties of a given phonological unit vary considerably depending on its context. The value of [e.g.,] lip aperture achieved, however, remains relatively invariant no matter what the vowel context (1995, p. 184).

The gesture is specified invariantly ... but the ... action is distributed across component articulators in a context-dependent way (1995, p. 185).

This latter feature -- invariance -- is a requirement if they are to serve as primitives of phonological theory. (Otherwise, there would be an open-ended set of primitives corresponding to the open-ended number of ways a particular gesture may be physically implemented.) The difficulty is that the units at the macroscopic-level descriptions -- the phonological gestural level -- are not actually invariant, or not necessarily invariant as a matter of empirical fact. Instead they are "defined" as such by the theory.

The values of the dynamical parameters associated with a ... gestures are macroscopic properties that define it as a phonological unit and allow it to contrast with other gestures... *These values are definitional*, and remain invariant as long as the gesture is active (1995, p. 184).

The difficulty here is not simply that to a certain extent the invariance of the phonological primitives is stipulative, but rather that it is problematic for the theory whether the invariance is "definitional" *or* empirical. If the gestural descriptions in terms of dynamic systems really are invariant as a matter of empirical fact, then they are well suited to serve as linguistic primitives (at least on that score). But to this extent they are not automatically isomorphic to their lower-level counterparts. For these are acknowledged to be variable. The problem is not the relative level of detail in the two descriptions, but rather that in one case, the units are interchangeable -- this token 'b' is as good as that token 'b' -- and in the other case they are not. So to the extent that the phonological gestures are invariant descriptions, it is not at all straightforward how these units solve the lack of invariance problem.

On the other hand, to the extent that we allow the dynamic systems that characterize gestures to vary as their lower level counterparts do, we have the complementary problem. Then the isomorphism between linguistic and physical gesture can be maintained, but at the linguistic level, they are not suitable primitives because they are not interchangeable. While the traditional approach to the lack of invariance problem flounders on the attempt to show that the physical segments really are invariant, from the point of view of the present criticism, articulatory phonology flounders on the attempt to show that the functional gestures really are invariant.

My aim thus far has not been to argue against the idea of gestures as phonological primitives *per se*. Rather, I've tried to show that in the effort to resolve the mismatch between speech viewed as a set of discrete linguistic units and speech viewed as a dynamic physical (production) process, the import of, and justification for, the linguistic-level description has been short-changed. As noted above, gestures as phonological primitives are said to be able to signal lexical contrast, but they are not primarily justified in terms of this role; nor is their role as such developed or treated as central ... as one would expect with a phonological theory. Granted different phonological theories have concentrated on different problem, each deals with a specifically linguistic domain or problem. Articulatory phonology on the other hand is primarily concerned to be able to describe phonological and physical levels as high- and low-level descriptions of the same event.

4. Critique of the Goal of “Unifying” Phonological and Physical Levels

As noted above, the methodological motivation for positing gestures as phonological primitives is clear: if the functional primitives of speech are defined in the same terms as the physical primitives, the traditional problem of how units at the two levels are related does not get off the ground, for the functional level is defined in such a way that its primitive elements are ipso facto transparently related to the physical primitives – i.e., the way gestures get to be functional primitives is just by being isomorphic to gestures at the articulatory, i.e., physical level. In this section I will argue that it is this motivating assumption -- that phonological primitives need to be isomorphic to physical primitives -- that is problematic. More specifically I shall argue that Browman and Goldstein are wrong to think that the only way phonetics and phonology can be systematically related is by being “unified” in the manner they describe.

Browman and Goldstein note quite rightly that functional and physical studies of speech have historically proceeded largely in independence from one another. Phonology studies the sound structure of language in functional terms. As a branch of theoretical linguistics it is related more immediately to other levels of linguistic structure such as morphology and syntax. By contrast the acoustic, auditory, and articulatory properties of speech are studied in phonetics, psychology, and speech science departments and are not especially constrained by results and changes in phonology. They are also right I think to see this indifference as a limitation. Nevertheless, I disagree with their basic approach to redressing this problem.

The main shortcoming of their approach to the problem of how the two levels are related is that they view the space of possibilities in too stark and narrow terms: either these two descriptions are “unified” or the two representations are thought to be totally “incommensurate”. Without worrying about the precise interpretations of this notoriously misunderstood term, we can identify other types of relations between these two kinds of descriptions that are neither unified (in the sense of invoking a single set of parameters) nor incommensurate in they sense that they there is no systematic way of describing the relation.

Two ways in which units at different levels may be systematically related without being “unified” in the above sense, are if the units at the two levels exhibit either *multiple realizability* or *context dependence*. Multiple realizability describes the case in which a single higher-level (functional) property is implemented by more than one kind of lower-level (physical) property. Context dependence describes the case in which a single lower-level (physical) property plays the role of different higher-level (functional) properties in different contexts. To adapt examples familiar from acoustic phonetics, the [d] sound plays the same functional role in the syllables [di] and [du] (i.e., it is heard as the same sound), but it is realized by different physical properties in the two cases. The primary acoustic cue for the [d] in [di] is thought to be a rising second-formant transition, while in the syllable [du] it is thought to be signaled by a falling transition. In this sense [d] may be thought of as multiply-realizable. Context dependence may be illustrated with the example of a single burst of noise centered on 1440 Hz which is heard as different higher-level properties (either a [p] or a [k]) in different phonetic contexts (when followed by an [i] or an [a] respectively).

The point here is not that multiple realizability and context dependence are sufficient to account to the relation between the two levels, but (i) they give us more explanatory options, and (ii) their existence shows that Browman and Goldstein can’t justify their approach simply by showing that the two levels are not “irrelevant” on their approach. One could agree that they should be related and mutually constraining without agreeing that this should be down in the way Browman and Goldstein suggest.

It’s not simply that they need to argue more for their particular solution; rather it’s that they are wrong to compromise the goals of phonological theory in order to achieve their solution. That they do so I have tried to argue above. That they don’t see this as a problem is shown by the fact that they criticize early attempts that focused on function and ignored the physical aspects of speech:

In general, what was seen as important about phonological units was their *function*, their ability to distinguish utterances (1995, p.178).

Well, as we have seen, and as Browman & Goldstein acknowledge elsewhere, that *is* what's important. Moreover, in identifying such functional units it is important to be able to largely ignore the physical and physiological details. Of course at the limits they provide a constraint, a unit which couldn't be pronounced or perceived effectively at all would not make a very good functional unit. But in their role of distinguishing utterances it is important and appropriate that these units not be tied to a particular implementation. After all, phonological units can be implemented not only with different e.g. aperture pressures, but with hands as well, as in ASL. The goal of a functional -- in this case, phonological theory -- is quite different from, and invokes units different from, a mechanical description. The former must be able to be implemented physically, but this does not mean it should be reducible to a physical description. As we have seen, a complex, expressive system such as language needs particulate structure. For this it needs discrete units. Natural phenomena aren't discrete; but there are good (particulate structure) reasons for analytic, linguistic, functional structure being this way. Discrete segmental structure is best viewed as an abstraction from the physical signal and processes that produce it (though it is no less real for being so). So one should focus on the phonological justification of phonological primitives, not because the physical level doesn't constrain at all, but because it doesn't do so by forcing us to posit either physical segments or functional gestures.

References

- Browman, Catherine P., and Goldstein, Louis (1990a). Gestural Specification using dynamically-defined articulatory structures. *Journal of Phonetics*, 18, pp. 299-320.
- Browman, Catherine P., and Goldstein, Louis (1990b). Representation and Reality: Physical systems and phonological structure. *Journal of Phonetics*, 18, pp. 411-424.
- Browman, Catherine P., and Goldstein, Louis (1992). Articulatory Phonology: An overview. *Phonetica*, 49, pp. 155-180).
- Browman, Catherine P., and Goldstein, Louis (1995). Dynamics and Articulatory Phonology. In Robert F. Port and Timothy Van Gelder (Eds.), *Mind as Motion: Explorations in the dynamics of cognition* (pp. 175-193). Cambridge, MA: The MIT Press.
- Liberman, Alvin M., and Ignatius G. Mattingly (1985). The Motor Theory of Speech Perception Revised. *Cognition*, 21, pp. 1-36.
- Studdert-Kennedy, Michael, and Goldstein, Louis (2003). Launching language: The gestural origin of discrete infinity. In M.H. Christiansen and S. Kirby, (Eds.), *Language Evolution: The States of the Art*. Oxford University Press.

Proceedings of the 2003 Texas Linguistics Society Conference: Coarticulation in Speech Production and Perception

edited by Augustine Agwuele,
Willis Warren, and Sang-Hoon Park

Cascadilla Proceedings Project Somerville, MA 2004

Copyright information

Proceedings of the 2003 Texas Linguistics Society Conference:
Coarticulation in Speech Production and Perception
© 2004 Cascadilla Proceedings Project, Somerville, MA. All rights reserved

ISBN 1-57473-402-4 library binding

A copyright notice for each paper is located at the bottom of the first page of the paper.
Reprints for course packs can be authorized by Cascadilla Proceedings Project.

Ordering information

Orders for the library binding edition are handled by Cascadilla Press.
To place an order, go to www.lingref.com or contact:

Cascadilla Press, P.O. Box 440355, Somerville, MA 02144, USA
phone: 1-617-776-2370, fax: 1-617-776-2271, e-mail: sales@cascadilla.com

Web access and citation information

This entire proceedings can also be viewed on the web at www.lingref.com. Each paper has a unique document # which can be added to citations to facilitate access. The document # should not replace the full citation.

This paper can be cited as:

Appelbaum, Irene. 2004. Physical Segments and Functional Gestures. In *Proceedings of the 2003 Texas Linguistics Society Conference*, ed. Augustine Agwuele et al., 1-8. Somerville, MA: Cascadilla Proceedings Project.

or:

Appelbaum, Irene. 2004. Physical Segments and Functional Gestures. In *Proceedings of the 2003 Texas Linguistics Society Conference*, ed. Augustine Agwuele et al., 1-8. Somerville, MA: Cascadilla Proceedings Project. www.lingref.com, document #1062.