Anim Cogn
DOI 10.1007/s10071-010-0328-7

ORIGINAL PAPER

Semantics of primate gestures: intentional meanings of orangutan gestures

Erica A. Cartmill · Richard W. Byrne

Received: 27 January 2010 / Revised: 14 April 2010 / Accepted: 18 May 2010

Abstract Great ape gesture has become a research topic of intense interest, because its intentionality and flexibility suggest strong parallels to human communication. Yet the fundamental question of whether an animal species' gestures carry specific meanings has hardly been addressed. We set out a systematic approach to studying intentional meaning in the gestural communication of non-humans and apply it to a sample of orangutan gestures. We propose that analysis of meaning should be limited to gestures for which (1) there is strong evidence for intentional production and (2) the recipient's final reaction matches the presumed goal of the signaller, as determined independently. This produces a set of "successful" instances of gesture use, which we describe as having goal–outcome matches. In this study, 28 orangutans in three European zoos were observed for 9 months. We distinguished 64 gestures on structural grounds, 40 of which had frequent goal–outcome matches and could therefore be analysed for meaning. These 40 gestures were used predictably to achieve one of 6 social goals: to initiate an affiliative interaction (contact, grooming, or play), request objects, share objects, instigate co-locomotion, cause the partner to move back, or stop an action. Twenty-nine of these gestures were used consistently with a single meaning. We tested our analysis of gesture meaning by examining what gesturers did when the response to their gesture did not match the gesture's meaning. Subsequent actions of the gesturer were consistent with our assignments of meaning to gestures. We suggest that, despite their contextual flexibility, orangutan gestures are made with the expectation of specific behavioural responses and thus have intentional meanings as well as functional consequences.

Keywords Great ape · Method · Meaning · Intention · Primate communication

Electronic supplementary material The online version of this article (doi:10.1007/s10071-010-0328-7) contains supplementary material, which is available to authorized users.

E. A. Cartmill · R. W. Byrne (✉)
Scottish Primate Research Group and Centre for Social Learning and Cognitive Evolution, School of Psychology,
University of St. Andrews, St. Andrews, Fife KY16 9JP, UK
e-mail: rwb@st-andrews.ac.uk

E. A. Cartmill
Department of Psychology, University of Chicago,
Chicago, IL 60615, USA

Introduction

The study of meaning in animal communication presents different challenges depending on whether one is interested in determining either the *message* (the information present in the signal[1]) or the *meaning* (as perceived by a recipient, based on the signal and the context) of apparently communicative acts (Smith 1965). In seeking to compare animal communication systems to human language, a third challenge presents itself: that of determining whether the signaller intends for the recipient to derive a particular meaning from the message contained in the signal. *Intentionality* is central to the comparison, because language relies upon an intentional structured transfer of information: actively selected and encoded by the speaker, and

[1] By *information present in a signal*, we mean any quality of a signal that is predictably associated with an external event or internal state of the signaller and is thus a reliable indicator of an entity, state, or act external to the production of the signal itself. For a discussion of the importance of defining information and meaning, see Rendall et al. (2009).

received and interpreted by the listener. Interpretation of an action as communication can vary dramatically depending on whether one approaches the event from the perspective of either the signaller or the receiver (Seyfarth and Cheney 2003). A signal may have meaning for a recipient, if the recipient is able to extract new information about the environment or the signaller, through a combination of the signal's message and the overall context in which it is produced; but the signaller might not have intended to convey that or any other meaning.

Where we lack evidence of the signaller's intent, the most that can be discerned is that some signals may have *functional meaning*; but where there is an indication that the signaller produced a signal with the goal of eliciting a particular behavioural outcome in the recipient, we can analyse it for *intentional meaning*. The functional meaning of a signal, inferred exclusively through recipient reaction, might even be different from the intentional meaning of that signal: for example, a smile in a bar might be intended as mere politeness by the signaller but perceived as a romantic advance by the recipient. Determining the intentions of a non-linguistic animal is especially difficult, since intentions are not directly observable, and it is not a straightforward task to detect intentions from observing behaviour. Nevertheless, researchers must somehow infer signallers' intentions through observable external clues, unless analyses are restricted to how a signal functions.

Studies of primate communication have focused almost entirely on the functional meaning of signals. This is because most research has been done on vocal communication, a medium that is ideal for investigating questions about the message and functional meaning of signals. The message of a call can be assessed by examining the acoustic parameters of vocalizations for reliable informational correlates: e.g. the callers' location, sex, mood, identity, etc. Functional meaning can be studied by using "playback", in which experimenters broadcast previously recorded calls to groups of animals and observe their reactions. Using this technique, experimenters can manipulate acoustic parameters to examine which variables are important in allowing listeners to extract specific meanings (e.g., Seyfarth et al. 1980a, b; Kroodsma 1986; Zuberbühler 2000; Arnold and Zuberbühler 2006). Although effective in investigating functional meaning, the playback method affords little progress towards understanding the intentional meaning of vocalisations because the methodology is based on listeners' reactions to calls.

Gestural communication may be more tractable for addressing questions of intentional meaning, at least when studying great apes. In great apes (hereafter, apes), gesture is thought to be more flexible than vocal communication (Call and Tomasello 2007a, b; Pollick and De Waal 2007), and evidence for intentional use of gestures has been found in several ape species (Pika et al. 2003, 2005a, b; Liebal et al. 2006; Call and Tomasello 2007a, b; Genty et al. 2009). Because at least some ape gestures are intentional actions, they provide a unique opportunity to study intentional meaning in primate communication. However, it is not yet possible to "play back" recorded gestures to apes in a naturalistic manner, and identification of acts as gestures can be problematic since the movements used in gestures overlap with those used in many other types of movement. To study intentional meaning in ape gestures, one must take into account the behaviour of both the signaller and the recipient.

Any movement an animal makes is capable of transmitting information about it to an observant bystander. It is possible to deduce that an animal is lame by observing its walk and to determine that it has parasites by observing its persistent scratching; yet, these potentially informative movements are primarily effective actions rather than communicative signals. Since gestures share a modality with many of the utilitarian movements of daily living, discerning which movements should be analysed as communication can depend to a great degree on identifying which movements are performed *as signals*. Researchers studying intentional behaviour typically require that an animal's actions are goal directed and show some degree of flexibility in reaching the goal: this decreases the likelihood of mistakenly labelling inflexible "hardwired" patterns of behaviour as *intentional*. Thus, the goal-plus-flexible-path combination features heavily in psychological literature from fields as diverse as child development (Piaget 1952; Bruner 1981), comparative cognition (Liebal et al. 2006; Pika 2008), and neuroscience (Dickinson and Balleine 1994).

If behaviour is not hard-wired, it may vary between individuals; thus, a movement that one individual uses as an intentionally communicative gesture may be used by another individual in a non-communicative way. In many previous studies (Pika et al. 2003, 2005a, b; Liebal et al. 2006; Call and Tomasello 2007a, b), it has been assumed that two observations of a particular gesture used in an intentional manner by any individual would suffice to conclude that the gesture was intentional whenever any individual used it. It may be unwise, however, to assume that a gesture is intentional for *all* individuals if there is real evidence only for intentional usage in one or a few individuals: individual differences may exist in the way apes communicate through different movements, as in humans. It is therefore desirable to establish intentional use of a movement *by a specific individual* before including that movement as a gesture in that individual's repertoire (see Genty et al. 2009).

Previous studies of ape gesture have revealed the flexibility of gesture use but have paid less attention to the

signallers' precise goals. All great apes have been found to use gestures flexibly in different contexts in captivity (Pika et al. 2003, 2005a, b; Liebal et al. 2006; Call and Tomasello 2007a, b; Pollick and de Waal 2007; Arbib et al. 2008; Genty et al. 2009). This *contextual flexibility* is characterized by the lack of a one-to-one correspondence between external events or social contexts and particular gestures, and in great apes it has been taken to imply social flexibility in achieving their aims: "disassociation of means and ends" (see Call and Tomasello 2007b). Certainly, great apes do show flexibility in their gestures; e.g. they choose effective signals in response to whether others can see them (Tanner and Byrne 1993; Liebal et al. 2004; Poss et al. 2006; Call and Tomasello 2007a, b) and continue to pay attention to the attentional state of their audience while signalling (Genty et al. 2009). However, gestures that appear in the same context may, in fact, have different goals: for instance, some gestures may be used to initiate play, some to refuse play requests, and some to end play in progress. Conversely, the same gesture may serve exactly the same goal in several contexts. An analysis of flexibility using context as a proxy for goal may therefore lead to spurious conclusions about the flexibility of certain gestures.

To date, only a handful of studies have attempted to identify the intentional meanings of natural ape gestures used with conspecifics. Although some exciting conclusions have been drawn, the methods used were not always explicit or systematic, and it is often difficult to compare results between studies. Call, Tomasello and colleagues (as reviewed in Call and Tomasello 2007a) conducted the only study of ape gestures that is fully comparable between species. This was a landmark study in that it used the same methods in several species to determine gestural repertoires and use. The authors catalogued the gestures of four species of great apes as well as siamangs (*Symphalangus syndactylus*), using multiple captive groups for each species. The study found more evidence of gestural flexibility than specificity; however, it was designed to compare gesture form and use between species and did not assess intentional meanings for specific gesture types. Studies focusing on a single gesture type or gestures of a single individual have concluded that some gestures have more specific meanings or uses, finding evidence of both deictic gestures (in captive and wild chimpanzees: Leavens et al. 2004, Pika and Mitani 2006) and iconic gestures (in a captive lowland gorilla: Tanner and Byrne 1996). These highly targeted studies suggest that apes' gestures have the potential to carry very specific intentional meanings. Only one analysis of meaning within an ape's entire gesture repertoire has been attempted thus far (Genty et al. 2009). In this study of the western gorilla, a gesture's meaning was estimated by the reactions of the target audience that apparently satisfied the gesturer as well as the stimulus conditions that elicited the gesture. Although gestures differed in their pattern of occurrence, all gestures were found to be "multi-purpose;" with an average individual using a gesture for 4.5 ± 1.7 functions. Only gestures that occurred singly or as the first gesture in a sequence were analysed for meaning. While this simplified the analysis by restricting it to a single signal and reaction in each case, it necessarily excluded some gestures from analysis.

We believe that a more systematic approach to gesture meaning needs to be adopted. It should be broadly comparable between species (Call and Tomasello 2007a) and based on the behaviour of the signaller (Genty et al. 2009) and yet allow extended interactions between individuals of a particular species to guide the analysis of meaning. Most importantly, any study of meaning in animal signals should allow the human attributions of signal meaning to be empirically tested. Though it is not possible to play back recorded gestures as one would recorded calls, the method we advocate here uses the signaller's behaviour following different recipient reactions to evaluate our previous attributions of gesture meaning.

Thus, we propose taking into account: (1) the apparent aim of the gesturing individual and (2) the reaction of the recipient that satisfies that goal. Where the two match consistently, across examples of a particular gesture, an intentional meaning is indicated for that gesture. We can then test that particular attribution by examining what the signaller does in cases where the recipient's reaction does not match the hypothesized meaning of the gesture used. Here, we apply this systematic and usage-based approach to intentional meaning to a sample of orangutan gestural communication. The aim is to identify gestures that are used predictably to elicit specific reactions and to begin to approach a lexicon of orangutan gestures and their intentional meanings.

We begin this analysis by identifying all orangutan movements that are made in social contexts and apparently achieve no primary function for the subject by their execution; we define these movements as *potential gestures*. Then, we use cases where instances of potential gestures are employed, in a recipient-directed and goal-oriented way, to determine which individuals use which movements intentionally as communicative *gestures*. Conservatively, where potential gestures are made by an individual that fails to demonstrate any clearly intentional use of that gesture type, we exclude those cases from further analysis.

To assess the intended meaning of a gesture, we begin by ascribing the apparent aim of the signaller (*presumed goal*) to each example of gesture, based on the individuals involved and the immediate social context in which the interaction occurs. We then filter out cases where the signaller does not appear satisfied by the outcome and restrict

attention to the remainder of cases, where we take it that signaller goal have been met by the outcomes of the interactions (*interaction outcome*). We define these cases of gesture use as having *goal–outcome matches*. The extent to which a single gesture type corresponds to a single goal–outcome match, for a particular individual, indicates the degree of precision or ambiguity in the *intentional meaning*. In attempting to validate these assignments of intentional meaning, we turn to signallers' behaviour, following recipient reactions that did not meet our assigned intentional meanings of the gestures used. We hypothesize that orangutans will continue to gesture following reactions that do not fulfil the request conveyed in the gesture. By identifying the intentional meanings of (some) orangutan gestures, we aim to make a first attempt at identifying the primary semantic contrasts present in orangutan gestural communication.

Methods

Subjects

Twenty-eight orangutans were observed in three European zoos. The sample was composed of two groups of Bornean orangutans (*Pongo pygmaeus*) housed at Twycross Zoo, UK, and Apenheul Primate Park, Netherlands, and one group of Sumatran orangutans (*Pongo abelii*) housed at the Durrell Wildlife Conservation Trust, Jersey. Following Rijksen (1978), we defined individuals under 2.5 years old as infants, those between 3 and 10 as juveniles (combining his subadult and juvenile age classes), and those 10 years and above as adults. These categories also largely parallel distinctions based on ecological competence in Sumatran orangutans (van Adrichem et al. 2006). The group at Twycross Zoo comprised 8 orangutans: 4 adults (1 M, 3 F), 2 juveniles (1 M, 1 F), and 2 infants (F). The group at Apenheul Primate Park comprised 12 individuals: 7 adults (2 M, 5 F), 3 juveniles (1 M, 2 F), and 2 infants (F). The collection at Durrell contained 8 orangutans: 4 adults (1 M, 3 F), one juvenile (M), and three infants (2 M, 1 F).

Procedure

We used ad libitum sampling to record video of social interactions we thought likely to involve gesture: active social interactions such as coordinated locomotion, playing, object manipulating, or foraging. This method was used because pilot study suggested that focal-animal sampling, such as that employed by Liebal et al. (2006) to record gesture in orangutans, would significantly reduce the number of gestures observed, since active social interaction occurred primarily in punctuated bursts. We observed the orangutans for between 5 and 7 h 5 days a week from areas accessible to the public. The animals were not filmed during periods of rest, solitary feeding, and solitary play.

Video was collected at Apenheul from March to April 2005 and November 2006, at Twycross in April, May, and October 2006, and at Durrell from June to August 2006. This schedule resulted in 3 months of observation at each of the three sites (exceeding 300 h of observation time per site), giving approximately 50 h of video. A Sony Handicam DCR TRV-38 was used to film the orangutans in both their indoor and outdoor enclosures.

Analysis

Identifying potential gestures

Movements of the face, head, limbs, or body were defined as *potential gestures* when they were both "motorically ineffective" (see Call and Tomasello 2007a, b) and directed towards another individual. Gestures were considered to be motorically ineffective if they did not directly perform a practical act such as scratching or picking up an object. Potential tactile gestures were considered to be motorically ineffective if the recipient did not move immediately (as if by force) after contact, or the degree of the actual movement was considered to be greater than justified by the force of the gesture. Two coders were used to check judgement of motoric-effectiveness. If the primary coder considered that a recipient acted as if moved by force or that a signaller used a potential gesture in a way that suggested the application of force, a second coder assessed the video and the potential gesture was discarded if either coder judged that the potential gesture was motorically effective. Determining whether a movement was directed towards another individual was straightforward when the movement was tactile. Visual movements were defined as directed only if the gesturing individual performed the act while oriented towards and apparently looking at another individual. We did not require that visual gestures be potentially detectable by the individuals towards whom they were directed (as did Genty et al. 2009), because this would have excluded cases of ineffective signal use and ruled out the possibility of analysing the frequency of ineffective signal use.

We coded all potential gestures that occurred during a social interaction unless the interaction involved continuous active contact between the individuals. Coding was thus stopped at the onset of physical play (such as wrestling), cuddling, nursing or carrying an infant, or mating. If continuous interaction ceased for at least 10 s, any gestures following the pause were again coded and included in the analysis. We included facial expressions in our list of

potential gestures, provided they met the criteria of directedness and were performed on their own rather than accompanying a manual potential gesture (when they accompanied a manual potential gesture, they were recorded as an aspect of that potential gesture).

Potential gestures were defined according to their structural similarities in the following variables: modality, body part, movement, force, speed, and use of an object. All variables were divided into mutually exclusive categories. For each observed example of a potential gesture, which we henceforth refer to as a "token" of that gesture, we coded 30 separate variables (see Table 1). These variables included the animals involved in the interaction, the specific form of the movement, the reaction of the recipient, and any subsequent gestures or reactions of the signaller.

Identifying gestures

In order to decide whether an individual used a potential gesture as an intentional signal (and therefore as a gesture, by our criteria), we required evidence that the individual used that potential gesture at least once in an intentional manner. To this end, we excluded from our dataset all potential gesture tokens produced by any signaller that failed to show any clearly intentional use of that same gesture. The final dataset used in our analysis therefore included only tokens of a particular gesture produced by signallers that had shown the ability to use that gesture intentionally at least once.

To determine whether tokens of potential gestures were used intentionally, we employed a widely used definition of intentionality in non-linguistic communication, in which an act is deemed to be intentionally communicative if it is: (1) directed towards another (part of our criteria for potential gestures), with (2) the apparent objective of obtaining a goal, and (3) employed flexibly rather than as an automatic response to a stimulus (Bruner 1981; Pika et al. 2005a; Tomasello and Call 2007). In order to establish whether the signaller had an intended goal, we looked for evidence that the signaller "expected" a reaction from the recipient, rather than signalling in an automatic response to some environmental stimulus; measures of expected reaction included remaining oriented towards the recipient without performing any additional act as if waiting for a reaction, alternating gaze between the recipient and an object or location, persisting, and using modalities appropriate to the attentional state of the recipient (e.g. using silent visual gestures only when the recipient is looking). Use of potential gestures in varying combinations with other potential gestures was also considered evidence of flexible, goal-directed behaviour.

We thereby rated every token of each potential gesture on the strength of evidence for intentional usage. A token was taken to *support* an intentional interpretation of that gesture, if it was directed towards another individual and the signaller indicated that a reaction was expected—by waiting for a reaction, looking towards the recipient's face, escalating the signal, or following the gesture with either another gesture or an effective action. A potential gesture was deemed to be a gesture for a particular individual only if he or she used it at least once in a manner that supported intentional use. Potential gestures were not included in an individual's repertoire if they failed to meet these criteria.

Identifying goal–outcome matches

The next step, in the determination of a gesture's intentional meaning, was to combine a human estimate of the signaller's intent (*presumed goal*) with an objective measure of the final outcome of the interaction between the two individuals (*interaction outcome*).

We ascribed a signaller goal to every gesture token in an interaction (Fig. 1b). We based our judgements of

Table 1 Variables coded for each potential gesture token

Social variable	Context within exchange	Structure of gesture	Resulting actions
Identity of gesturer	Did the gesturer initiate the interaction?	Gesture	Presence of response waiting
Identity of recipient	Total number of gestures in exchange (both parties)	Modality	Reaction
Gesturer's age class	Position of gesture in exchange	Facial expression	Presence of persistence
Relationship between gesturer & recipient	Number of gestures in sequence (gesturer's gestures)	Body part used	Type of persistence
Context	Position of gesture in sequence	Handshape (if applicable)	2nd gesture
Gesturer's visual attention		Object (if used)	2nd gesture limb used
Recipient's visual attention		Target location	3rd gesture
Recipient's behaviour before gesture			Fourth gesture
Presence of observers not involved in exchange			Outcome of interaction

presumed goals on (1) the general context, including the recent course of interaction between the participants, (2) our knowledge of the relationship between the signaller and recipient (e.g. subordinate to dominant adult vs. infant to mother), as well as on (3) the form of the gesture itself and the result it seems designed to effect (e.g. gestures made towards the body are more likely to encourage approach than departure). The aim was to ascribe signaller goals to each gesture token, independent of the immediate reactions of the recipient or the interaction outcome of that exchange. This meant that we could be surprised when an interaction outcome that did not meet our presumed goal caused the signaller to stop gesturing. We also did not assume that the goals of every gesture within a sequence would be the same. Thus, an interaction outcome might match the ascribed goals of some gestures within a sequence but not others. We gave each gesture one of eight presumed goals: *Affiliate/Play*, *Stop action*, *Sexual contact*, *Look towards signaller, Look at/Take object*, *Share food/object*, *Co-locomote*, or *Move away*. Inter-observer reliability of two experimenters' attributions of gesturers' presumed goals was performed on 6% of the total dataset (90 gesture tokens) and achieved a high level of agreement ($k = .83$). Presumed goal was deemed therefore to be a replicable measure and a useful tool in further analysis of gesture intentional meaning.

For every social interaction involving one or more gestures, we recorded the immediate *reactions* of the recipient to each gesture. Reactions were categorized as: moving towards or away, sharing or taking food or objects, following or picking up the other, looking towards the other, playing, affiliating (including nursing and grooming), mating, ceasing an action, performing a particular gesture, or not responding at all. We defined the final reaction of the recipient as the *interaction outcome*, as it caused the signaller to stop gesturing (Fig. 1a). In interactions that consisted of a single gesture, the reaction immediately following that gesture was the interaction outcome. In longer interactions, the final reaction of the recipient was considered the interaction outcome for all the signaller's gestures throughout the interaction. We used the interaction outcome as the outcome for all gestures in a sequence, rather than the recipient's immediate reactions to each, because it was the only reaction that caused the signaller to stop gesturing and thus had the highest likelihood of fulfilling the goal of at least one gesture in the sequence.

We defined all gesture tokens whose presumed goals matched the interaction outcomes as having *goal–outcome matches* (Fig. 1c). In a goal–outcome match, the outcome of the interaction had to match the presumed goal of the gesture exactly, except for cases when the only presumed goal we were able to ascribe was to gain the attention of the recipient. (In that case, all affiliative outcomes were counted as matching even if they included other actions, because recipients always attended to the signaller before engaging in an affiliative interaction.)

Assigning intentional meaning

To investigate intentional meanings, we narrowed our dataset to gesture tokens defined as having goal–outcome matches. Gestures that were used predictably with a single goal–outcome match were considered to have a specific intentional meaning (corresponding to the signaller's presumed goal). We defined several levels of intentional meaning depending on how often a gesture was found to have a particular goal–outcome match. Gestures that were used with essentially a single goal–outcome match, i.e. more than 70% of the time and three times as often as the next most common goal–outcome match, were defined as having "tight meanings." Gestures used frequently with a single goal–outcome match, i.e. 50–70% of the time and twice as often as the next most common goal–outcome match, were defined as having "loose meanings." We defined gestures that had a single goal–outcome match less than 50% of the time as having "ambiguous meanings."

We then used the attributions of meaning to test our initial gesture definitions: that is, we explored whether gestures judged as having loose or ambiguous meanings could be redefined to have greater meaning specificity if combined with other gestures or split into more than one gesture using an additional structural variable in their definitions. We reasoned that meaningful gestures were likely a better reflection of the gestural distinctions made by the orangutans themselves, whereas our own categorization was fallible. By reconsidering the initial definitions of gestures, based on whether they demonstrated consistent goal–outcome matches, we were able to test our assumptions about where gesture boundaries lay, as well as to consider the granularity of our definitions and explore which variables were important in determining one meaningful gesture from another (Cartmill and Byrne 2010).

Testing assignments of intentional meaning

To evaluate the success of this exercise, we took the intentional meanings, derived from the subset of gesture tokens with goal–outcome matches, and applied them to all tokens of those gestures in the larger dataset. For this, both tight and loose gesture meanings were used. We compared the rate of a signaller's persisting, when the recipient's immediate reaction matched the gesture's intentional meaning, to their rate of persisting when the recipient's immediate reaction did not match the gesture's intentional meaning. Greater persistence when a signaller's apparent

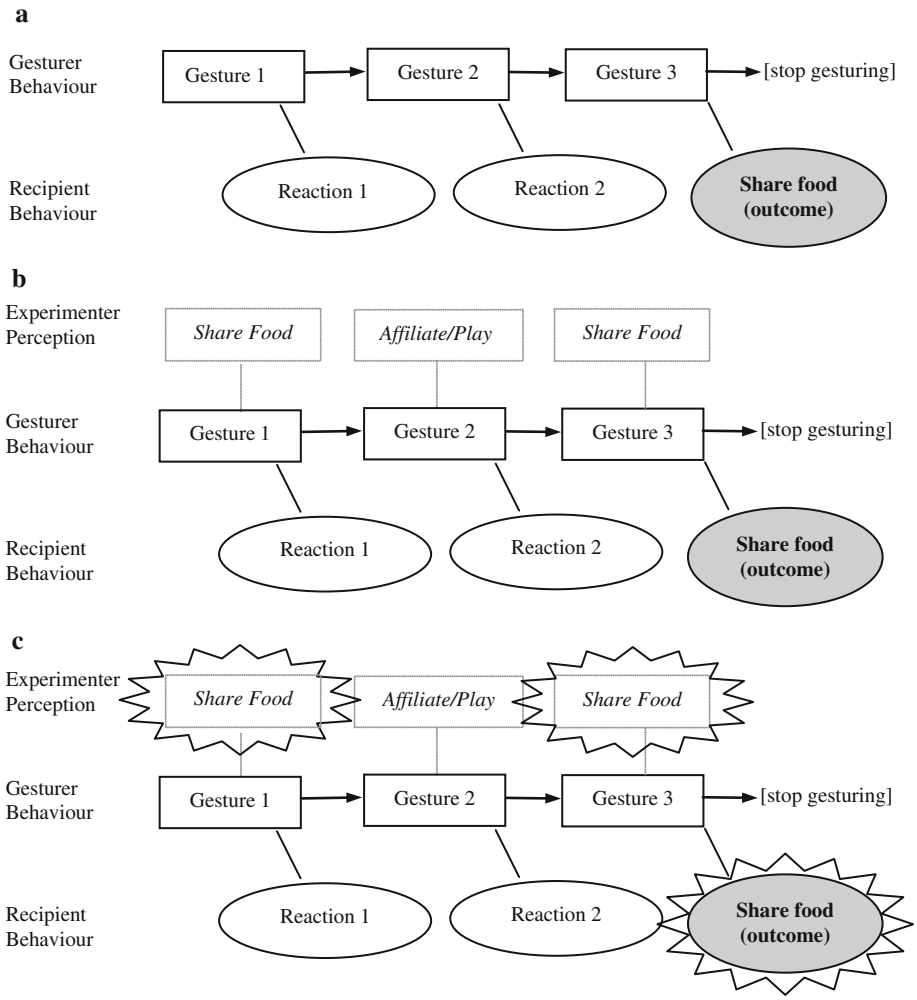

Fig. 1 Assigning meaning to gestures. The diagrams show the three stages in determining the meanings of gestures, based on observable events (gestures and reactions) and based on goal–outcome matches. **a** Directly observable gestures and reactions in a sequence of gestures. **b** Gestures, reactions, and experimenter-ascribed goals of the gesturer in a sequence of gestures. **c** Goal–outcome matches in a sequence of gestures. Note that both Gesture 1 and Gesture 3 have goal–outcome matches

aim was not realized would validate our provisional assignment of that intentional meaning.

Results

Goal–outcome matches

Applying our criteria for intentional use reduced the initial corpus of 1581 potential gesture tokens to 1,344 gesture tokens. Those 1,344 tokens were then categorized by their structural similarities into 62 gesture types. A table of all gesture types, with descriptions and token counts, is included as electronic supplementary material. When we further restricted attention to those gesture tokens with goal–outcome matches, the dataset reduced to 698 gesture tokens, including at least one token of each of the 62 gesture types. Of the 646 gesture tokens whose outcomes did not explicitly match their presumed goals, 344 received no reaction and 99 resulted in rejection (that is, the recipient turned away or pushed the signaller away). Two hundred and three gesture tokens (15% of the total gestures recorded) ended in active interaction outcomes that did not apparently match their presumed goals.

Tight, loose, and ambiguous meanings

In order to determine whether any gestures had intentional meanings, that is, were used reliably with particular goal–outcome matches, we focused on those gestures that were used at least four times in conjunction with any particular goal–outcome match. We imposed this minimum threshold of gesture use to ensure that our attributions of meaning would be based on several cases of use, particularly since we aimed to estimate the specificity of a gesture's meaning from the comparative frequency of its use for different purposes. Twenty-four of the 62 gestures were used fewer than four times towards any one goal and were thus excluded from detailed analysis of intentional meaning, leaving 38 gestures to be analysed.

The degree of specificity of intentional meaning was estimated by the probability that gestures would be used with

particular goal–outcome matches. Twenty-seven gestures met the criteria for tight meaning (produced more than 70% of the time with one goal–outcome match and three times as often as with the next most common goal–outcome match). Eight gestures met the criteria for loose meaning (produced between 50 and 70% of the time with one goal–outcome match and twice as often as with the next most common goal–outcome match). Three gestures were considered to have ambiguous meaning (produced less than 50% of the time with one goal–outcome match). To determine whether apparent ambiguity of any gesture resulted from *some* individuals having different meanings for the gesture, we checked whether the second most common goal–outcome match for each gesture was the product of a few individuals using the gesture exclusively with a different intentional meaning; this was not the case. Some individuals did appear to use gestures more ambiguously than other individuals, but we did not have a large enough dataset to investigate tightness of meaning on an individual level.

By reconsidering the definitions of gestures with loose or ambiguous meanings, we found that we could create two more gestures with tight meanings by grouping tokens of several ambiguous gestures according to the target location of the gesture (the area to which the gesture was directed). Once we added these two new gestures (*seize* and *pull away appendage*), we ended up with 64 total gestures; 29 with tight meanings, 7 with loose meanings, 4 with ambiguous meanings, and 24 that were used too infrequently to be analysed.

We found that that orangutans used specific gestures to communicate six of the eight intentional meanings that we considered originally: *Affiliate/Play*, *Stop action*, *Look at/Take object*, *Share food/object*, *Co-locomote*, and *Move away* (Table 2). Of these, *Affiliate/Play* was by far the most common: 25 gestures were found to have a tight or loose meaning of *Affiliate/Play*. The intentional meanings *Sexual contact* and *Look towards signaller* were not conveyed consistently by any gesture. Neither the forms nor the meanings of gestures appeared to vary systematically between different captive groups, suggesting that those gestures produced frequently enough to be included in our analysis were not group-specific in either their forms or use (see Table 2). Only one gesture, *fake*, was limited to use in one group, and it was idiosyncratic to a single adult female and observed only during a particular teasing exchange with her juvenile daughter. It is possible that this gesture was unique to the female or it could be that the context was so specific it was never observed in another group.

Validating assigned intentional meanings

To determine whether our ascription of intentional meanings to gestures had accurately captured a signaller's intent, we examined whether signallers were more likely to persist following immediate recipient reactions that failed to match the intentional meaning of the gesture used. Persistence was defined as repeating the same gesture, using another gesture, or holding the final position of the gesture for more than 2 s. Using gestures to which we had attributed intentional meanings, we analysed the subset of tokens that received active recipient reactions ($N = 626$), thus excluding all tokens where the recipient failed to react at all. When the recipient's reaction matched the assigned gesture meaning, signallers persisted 22 times and ceased gesturing 150 times. When the recipient's reaction did not match the assigned gesture meaning, signallers persisted 215 times and ceased gesturing 239 times. Thus, whether the recipient reaction matched the ascribed meaning significantly affected the signaller's probability of persisting ($\chi^2 = 63.35$, $df = 1$, $P < 0.001$).

Discussion

Our analysis of the intentional meanings of gestures incorporated both the apparent goals of signallers and the final outcomes of interactions into a single measure of goal–outcome matching. This approach highlighted the consistency in intentional meaning across contexts and led us to conclude that at least 29 orangutan gestures have "tight" meanings. When successful in communicative effect, these gestures were associated with a single intentional meaning more than 70% of the time. A further 7 gestures communicated a single intentional meaning more than half the time, whereas 4 appeared ambiguous between two or more meanings. The correlation between gesture form and meaning and the orangutans' persistence following responses that did not match their gestures' meanings demonstrate the effectiveness of the proposed method as a way of identifying gesture meanings. It is important to note, however, that 24 of the 64 gestures identified could not be analysed using this method because they occurred too rarely with goal–outcome matches. Our method of determining gesture meanings is a robust tool for large datasets but cannot be reliably used on infrequent gestures or small datasets.

We found that six intentional meanings were predictably associated with specific gestures. These six meanings—*Affiliate/Play*, *Stop action*, *Look at/Take object*, *Share food/object*, *Co-locomote*, and *Move away*—do not begin to approximate the highly referential meanings of human words but seem to represent a limited set of imperative requests helping to moderate dyadic interaction. This interpretation supports a growing body of work suggesting that ape gesture should be viewed as a dynamic interaction between two (or more) individuals rather than as

Table 2 Intentional gestures grouped by tight and loose meanings

Goal–outcome match	Tight meanings (70% and 3 times rate of secondary goal)				Loose meanings (50–69% and twice rate of secondary goal)			
		A	D	T		A	D	T
Affiliate/play	Air bite	5	1	8	Grab	27	15	26
	Arms up	5	0	12	Reach	16	13	20
	Back roll	1	2	5	Swat	13	9	15
	Bite	22	7	9	Touch	53	33	34
	Dangle	22	6	1				
	Duck lips open	0	11	0				
	Embrace	6	2	13				
	Fake	0	0	9				
	Hit	33	10	25				
	Hit ground/object	21	4	10				
	Play face	0	2	3				
	Pull hair	21	5	10				
	Put object on head	3	2	2				
	Raise arm	7	2	9				
	Raspberry face	4	0	1				
	Roll on back	2	5	7				
	Shake object	6	1	2				
	Simultaneous hit	2	0	7				
	Somersault	0	2	40				
	Swing	4	4	3				
	Wave	19	8	6				
Move away	Nudge	14	11	9	Push	15	14	29
	Shoo	6	8	5				
Share food/object	Food beg orally	5	4	4	Mouth	0	7	4
	Seize	8	16	19				
Stop action	Pull away appendage	5	2	3	Brush	10	3	12
Co-locomote	Tandem walk	0	0	18				
	Embrace pull	4	2	1				
Look at food/object	Offer	2	6	1				

The 29 gestures with tight meanings were defined using the 320 gesture tokens with goal–outcome matches, and the 7 gestures with loose meanings were defined using the 204 tokens with goal–outcome matches. The total number of tokens of each gesture observed at Apenheul (A), Durrell (D), and Twycross (T) is indicated to the right of the gesture

communicating referential information in the manner of many human words (King 2004; Genty and Byrne 2009).

One would expect gestures to be better suited for moderating dyadic social interaction than for sharing information broadly as they cannot be broadcast to many individuals at once. The perception of gestures generally relies on visual or physical contact, a highly limiting factor in communicating with multiple individuals. Signals conveying information about predators would be more effectively disseminated acoustically, as they can reach a greater number of individuals at a time and do so indiscriminately. Similarly, acoustic signals are ideal for communicating the signaller's location to other members of a group since the calling individual may be out of contact or far away from some group members. Simply, by producing an acoustic signal, an animal provides information about its location that can be used by other members of its group to maintain group cohesion, used by other members of its species to locate potential mates or competitors, or exploited by potential predators (Rendall et al. 2009). In contrast to signals broadcast to many other individuals, there might be an advantage to being able to conduct dyad-specific interactions in a more discreet medium like gesture, particularly in a primate group where interactions between lower ranking individuals may be monitored or interrupted by dominant ones.

If orangutan gestures function as dyadic social moderators, it may be useful to think of their intentional meanings as modifications of positive or negative intent to interact with others. The six intentional meanings we found can be grouped into positively and negatively valenced categories of requests. The positive requests (*Affiliate/Play*, *Co-locomote*, *Look at/Take object*, and *Share food/object*) fall into three classes of desires concerning (1) social affiliation, (2) co-locomotion, and (3) object transfer. We suggest that a fourth class of positive request, "sexual interaction", may be part of this system under natural conditions, but because our dataset did not include enough successful requests for sexual contact, we could not determine whether any gestures had this meaning. The requests *Move away* and *Stop action* can be grouped together into a single negatively valenced class of desire concerning disengaging or refusing interaction with others. Figure 2 illustrates our proposed classification of the intentional meanings of the orangutan gestures we recorded; we hope that this model will aid comparison of gesture meaning between different age classes, populations, or species.

Our conclusion that orangutan gestures carry specific intentional meanings seems to contrast that of Liebal et al. (2006), but the results may not be incompatible. Liebal et al. found that 81.8% ($N = 36$) of the gestures they observed were used in more than one social context and concluded that gestures were used flexibly across contexts; the frequency of use in each context was not given. Our findings suggest a tighter relationship between gestures and goals, but our analysis was based on the frequency that a gesture occurred with a particular social goal rather than on identifying which gestures occurred in more than one context. A one-to-one correspondence between goal and gesture is not necessary for a gesture to be described as having a particular intentional meaning; in fact, a complete lack of variation in use might imply that the behaviour was an unintentional innate response to a particular stimulus. A gesture with an intentional meaning should be expected to allow occasional variation in the correspondence between goal and outcome. This is already accepted in studies of primate vocalizations, where calls are not required *always* to be given in the expected contexts to be labelled as functionally meaningful (e.g. Seyfarth et al. 1980a, b; Zuberbühler 2001).

Without the ability to measure functional meaning through playback studies (as is done in vocalization research), gesture researchers must attempt to incorporate both production and reaction into their analyses and address the question of intentional meaning. The level of specificity we were able to identify in the intentional meanings of orangutan gestures resulted from our novel use of goal–outcome matching as a means of incorporating signaller intentions into the analysis of signal meaning. When paired with a dataset selected for signals demonstrating a high frequency of intentional use, goal–outcome matching is a strong tool for identifying intentional meaning.

The intentional meanings that we ascribed to gestures, based on their consistent use with particular goal–outcome matches, were generally validated by orangutans' increased likelihood of persistence following reactions that did not correspond to the intentional meaning of the gesture. The fact that orangutans sometimes persisted following a desired reaction may have been a product of the breadth of the categories of intentional meaning assigned to both the gestures and reactions. For example, all types of affiliation and play were combined into a single category that included actions as diverse as chasing, cuddling, wrestling, playing with objects, and sitting in contact. It is perhaps likely that the cases where the signaller continued to gesture, following a reaction that we recorded as fulfilling the intentional meaning *Affiliate/Play*, represent cases where the recipient reacted in a way that was broadly correct but that did not meet the precise expectations of the signaller (Cartmill and Byrne 2007).

Attributions of intentional meaning to gestures though the use of goal–outcome matching may be biased in favour of the meanings used by individuals who are "successful" communicators: determining intentional meaning in this way depends upon using those interactions in which the signaller is ultimately successful in achieving his or her goal. The likelihood of gesture success is greatly affected by the age class of the gesturer in relation to the recipient. Immatures (infants and juveniles) gesturing to adults receive no reaction 28% of the time, compared to 13% in the case of adults gesturing to immatures. Immatures thus "fail" in their gestures more than twice as often as adults, so the gestures of adults may be overrepresented in an analysis of intentional meaning based on goal–outcome matching.

Our results demonstrate that the gestural communication system of orangutans is composed of both ambiguous and intentionally meaningful gestures. Use of goal–outcome matching proved effective in incorporating both the goal of the gesturer and the final reaction of the recipient into a single measure of intentional rather than merely functional meaning. The frequency-based approach to intentional meaning revealed that orangutans use some gestures predictably with specific intentional meanings. As indicated in Fig. 2, these meanings reflect a semantic system with remarkably few contrasts compared to human language. Though more than half of the orangutan gestures we were able to analyse had predictable intentional meanings, they were still used flexibly in relation to the gaze or knowledge state of the recipient (Liebal et al. 2004; Poss et al. 2006; Cartmill and Byrne 2007). Flexibility and meaning need not be mutually exclusive, just as a signal can be both intentional and species typical (see Genty et al. 2009).

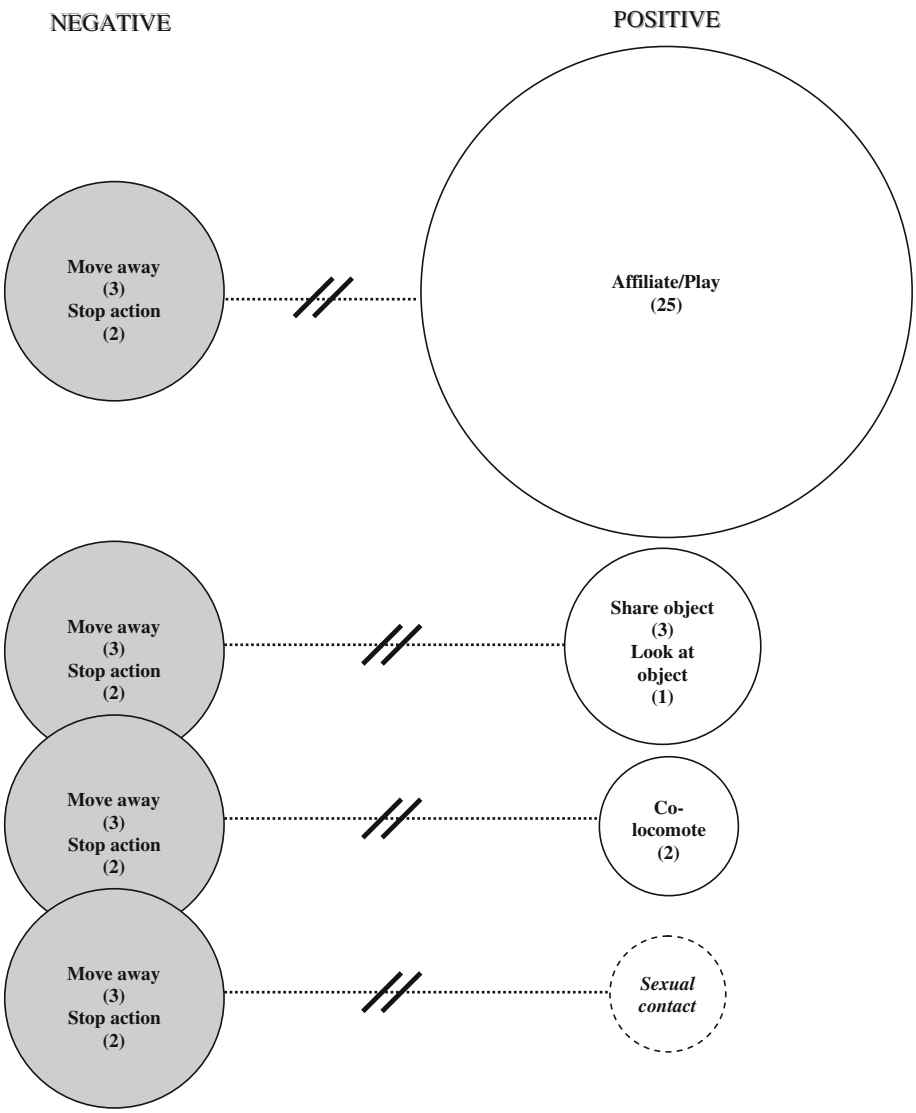

Fig. 2 Semantic contrasts evident in orangutan gestures. The terms in the *unshaded circles* are positive, interactive goals and have been clustered into three different intentional meaning groups based on contextual similarities. The three groups comprise gestures used to achieve social affiliation, object transfer, and co-locomotion, respectively. The terms in the shaded circles represent the contrast to each of these goals, a single intentional meaning group of rejection or disaffiliation. This group is repeated in the figure, shown as opposite to each positive gesture group, since the gestures it represents could be used to negate any of the positive requests. The relative sizes of the circles are proportional to the number of different gestures that fall within the group. The number of gestures associated with each meaning (through either a tight or loose association) is presented in brackets underneath the meaning. We include the presumed goal "sexual contact" because certain gestures were only used in conjunction with this goal; there were not, however, enough examples of goal–outcome matches for those gestures to be included in the analysis of meaning

Studying communication from the perspective of both the signaller and recipient provides a more complete picture of the way in which an action can be used as an intentionally meaningful signal. Our approach to studying intentional meaning in orangutan gestures begins this process and opens the door to a more cognitive approach to non-human communication.

Acknowledgments The authors would like to thank the staff and directors of Apenheul Primate Park, Durrell Wildlife Conservation Trust, and Twycross Zoo for their support and encouragement of this study. We warmly thank J. G. Foster, M. Cartmill, and K. Brown for their reviews of the manuscript, as well as J. C. Gomez and P. Lee for their comments on the methodology. Funding for the project came from the University of St Andrews and the Russell Trust.

Conflict of interest statement The authors declare that they have no conflict of interest with any of the organizations involved in the study.

References

Arbib MA, Liebal K, Pika S (2008) Primate vocalization, gesture, and the evolution of human language. Curr Anthropol 49:1053–1076

Arnold K, Zuberbühler K (2006) Semantic combinations in primate calls. Nature 441:303

Bruner J (1981) Intention in the structure of action and interaction. In: Lipsitt L (ed) Advances in infancy research. Ablex, Norwood, pp 41–56

Call J, Tomasello M (2007a) The gestural communication of apes and monkeys. Lawrence Erlbaum Associates, Inc, Mahwah

Call J, Tomasello M (2007b) The gestural repertoire of chimpanzees (*Pan troglodytes*). In: Call J, Tomasello M (eds) The gestural communication of apes and monkeys. Lawrence Erlbaum Associates, Inc, Mahwah

Cartmill EA, Byrne RW (2007) Orangutans modify their gestural signaling according to their audience's comprehension. Curr Biol 17:1–4

Cartmill EA, Byrne RW (2010) Addressing the problems of intentionality and granularity in non-human primate gesture.

In: Stam G (ed) Integrating gestures: the interdisciplinary nature of gesture. John Benjamins, Amsterdam (in press)
Dickinson A, Balleine B (1994) Motivational control of goal-directed action. Learn Behav 22:1–18
Genty E, Byrne R (2009) Why do gorillas make sequences of gestures? Anim Cognit 13:287–301. doi:10.1007/s10071-009-0266-4
Genty E, Breuer T, Hobaiter C, Byrne RW (2009) Gestural communication of the gorilla (*Gorilla gorilla*): repertoire, intentionality and possible origins. Anim Cognit 12:527–546
King BJ (2004) The dynamic dance. Nonvocal communication in African great apes. Harvard University Press, Cambridge
Kroodsma DE (1986) Design of song playback experiments. Auk 103:640–642
Leavens DA, Hopkins WD, Thomas RK (2004) Referential communication by chimpanzees (*Pan troglodytes*). J Comp Psychol 118:48–57
Liebal K, Pika S, Call J, Tomasello M (2004) Great ape communicators move in front of recipients before producing visual gestures. Interact Stud 5:199–219
Liebal K, Pika S, Tomasello M (2006) Gestural communication of orangutans (*Pongo pygmaeus*). Gesture 6:1–38
Piaget J (1952) The origins of intelligence in children. Norton, New York
Pika S (2008) Evidence for intentional and referential communication in great apes? In: Zlatev J, Racine T, Sinha C, Itkonen E (eds) The shared mind: perspectives on intersubjectivity. John Benjamins Publishing Company, Amsterdam, pp 165–186
Pika S, Mitani J (2006) Referential gestural communication in wild chimpanzees (*Pan troglodytes*). Curr Biol 16:191–192
Pika S, Liebal K, Tomasello M (2003) Gestural communication in young gorillas (*Gorilla gorilla*): Gestural repertoire, learning and use. Am J Primatol 60:95–111
Pika S, Liebal K, Call J, Tomasello M (2005a) The gestural communication of apes. Gesture 5:41–56
Pika S, Liebal K, Tomasello M (2005b) Gestural communication in subadult bonobos (*Pan paniscus*): Repertoire and use. Am J Primatol 65:39–61
Pollick AS, de Waal FBM (2007) Ape gestures and language evolution. Proc Natl Acad Sci USA 104:8184–8189
Poss SR, Kuhar C, Stoinski TS, Hopkins WD (2006) Differential use of attentional and visual communicative signaling by orangutans (*Pongo pygmaeus*) and gorillas (*Gorilla gorilla*) in response to the attentional status of a human. Am J Primatol 68:978–992
Rendall D, Owren MJ, Ryan MJ (2009) What do animal signals mean? Anim Behav 78:233–240
Rijksen HD (1978) A field study on Sumatran orangutans. Mededelingen Landbouwhogeschool, Wageningen
Seyfarth RM, Cheney DL (2003) Signalers and receivers in animal communication. Annu Rev Psychol 54:145–173
Seyfarth RM, Cheney DL, Marler P (1980a) Monkey responses to three different alarm calls: evidence for predator classification and semantic communication. Science 210:801–803
Seyfarth RM, Cheney DL, Marler P (1980b) Vervet monkey alarm calls: semantic communication in a free-ranging primate. Anim Behav 28:1070–1094
Smith WJ (1965) Message, meaning and context in ethology. Am Nat 99:405–409
Tanner JE, Byrne RW (1993) Concealing facial evidence of mood: perspective-taking in a captive gorilla? Primates 34:451–457
Tanner JE, Byrne RW (1996) Representation of action through iconic gesture in a captive lowland gorilla. Curr Anthropol 37:162–173
Tomasello M, Call J (2007) Introduction: Intentional communication in nonhuman primates. In: Call J, Tomasello M (eds) The gestural communication of apes and monkeys. Lawrence Erlbaum Associates, Inc, Mahwah, pp 1–15
van Adrichem GGJ, Utami SS, Wich SA, van Hooff J, Sterck EHM (2006) The development of wild immature Sumatran orangutans (*Pongo abelii*) at ketambe. Primates 47:300–309
Zuberbühler K (2000) Referential labelling in Diana monkeys. Anim Behav 59:917–927
Zuberbühler K (2001) Predator-specific alarm calls in Campbell's monkeys, *Cercopithecus campbelli*. Behav Ecol Sociobiol 50:414–422