Managing the Multiplicity of Meaning

Jacob G. Foster
University of California, Los Angeles, USA
foster@soc.ucla.edu

Erica A. Cartmill
University of California, Los Angeles, USA
cartmill@anthro.ucla.edu

Meanings multiply in Chuck Goodwin's lab. If anything characterizes his inimitable professional vision (Goodwin 1994), it is this: under Chuck's gaze, the possibilities pregnant in a piece of data multiply gloriously and generatively. Start with a video of a young girl cooking with her mother, and you may end up talking about the role of play in human evolution. Chuck's brain is without disciplinary borders; his lab is a social manifestation of that stimulating mental world. Through his special alchemy, diverse ideas and traditions join together in surprising transmutations. Although Chuck is far too modest to say so, the great work we do together in his lab plumbs deep questions in philosophical anthropology (Goodwin 2018).

Chuck's lab has been a crucible in which the two of us can explore our converging scholarly interests: Erica's in the evolution and development of human language, and Jacob's in the interplay between culture, cognition, computation, and action. In this essay, we sketch some thoughts, first incubated in Chuck's lab, that struggle with a fundamental puzzle in human communication: the fact that meaning can multiply without bound. How on earth do humans manage this multiplicity? We use the term "essay" in its original sense of an "attempt;" the treatment here is neither systematic nor supported by a detailed bibliography. In the spirit of Chuck's lab, we wanted to use our Festschrift contribution to play around with these ideas, which we hope to develop (and locate more thoroughly in the literature) in a longer paper. We think it's only fitting that this is the first academic paper we've written together; we have, in Chuck and Candy, an incredible model of two scholars, with distinct and powerful ideas and ways of thinking about the world, who create intellectual fireworks together – on the page and in the seminar room.

As an entrée to multiplying meaning, consider, first, the gestural communication system of great apes. Although their gestural lexicon is relatively large (e.g.,

~64 distinct gestures in orangutans), the range of possible meanings conveyed by those gestures is quite small (Cartmill, Byrne 2010; Byrne *et al.* 2017)[1]. In contrast to this semantic poverty, human communication is positively extravagant; indeed, our space of possible meanings is infinite and ever expanding! Building on the work of Sperber and Wilson (1996), Thom Scott-Phillips (2015) argues that the semantic richness of human language has a dual origin: (1) the linguistic code (conventional, learned maps from symbol to meaning plus a powerful syntax) and (2) our combined capacity for ostension ("the expression and recognition of communicative… intentions") and inference (in this case, the imputation of meaning from evidence provided by both the speaker and the environment)[2]. Scott-Phillips contrasts this system with the "natural code" of great ape gestures and vocalizations, which makes little use of ostension and inference (O-I) on his account. While we think Scott-Phillips has understated the O-I capacities of our great ape relatives (see Cartmill's commentary on Scott-Phillips 2015), we broadly agree with his account of human communication. We disagree, however, with his claim that "the linguistic code makes [the O-I system] expressively powerful." Or rather, that it *only* contributes to the O-I system's expressive power. Instead, we argue that a conventional linguistic code helps to manage the infinity of meaning that is *already possible* once a mature ostensive-inferential system is present.

To be clear, this explosion of meaning encompasses more than the familiar "gavagai" problem that confronts a language learner (Quine 1960). It extends beyond the well-worn problem of linguistic ambiguity (Levine 1988) and the fact that "linguistic meaning underdetermines speaker meaning" (Scott-Phillips 2015). In fact, the issue we raise here goes beyond conventional linguistic systems entirely, sharing a certain family resemblance with the famous "frame problem" in artificial intelligence, which notes that the "mind's central processes… can draw on information from any source," that "anything could be relevant" (Shanahan 2016). To put it bluntly: Ostension allows *anything* to be a meaningful, communicative act; not just words, not just gestures, but anything, from the tilt of a cap to a cryptic mark on a tree. The meaning that should be inferred from that communicative act is, in principle, completely arbitrary and independent of the features of the act. The same word, the same gesture, the same tilt, the same glyph, can mean one thing – or its opposite (see Scott-Phillips 2015 for amusing examples).

We can put this a bit more formally. The gestural communication system of great apes, as currently understood, is a (roughly) many-to-one map from the modestly sized space of signals to the very small space of meanings. In other words, many distinct signals map onto the same meaning (for example, the

[1] As Cartmill (2016) has argued, this finding must be taken with a substantial grain of methodological salt: scholars of great ape communication have typically looked for simple and reliable maps from signals to meanings, and their methods may have overlooked more sophisticated meaning-making of the type we describe in humans.

[2] There is an interesting parallel with Tomasello's (2010) distinction between linguistic code and common ground as joint foundations of meaning.

orangutan gestures *raise arm*, *air bite*, and *hit ground* are all used to initiate play). Although a given signal may map onto multiple meanings, the recipients of a signal nevertheless face a relatively easy task: because there is only a small range of likely meanings (sometimes a single meaning) for a given signal, recipients just need to learn and apply the signal-meaning maps. The human communication system, by contrast – thanks to ostension and inference – is radically one-to-many: a single signal may map onto many possible meanings[3]. Therein lies the problem: the signal, alone, cannot determine which of the many possible meanings applies in this situation. In principle, the set of *potential* meanings corresponding to a given signal is vast. This poses significant inferential problems for the recipient (see Foster *under review* for a detailed analysis of such computational questions in cultural learning and cognition).

Chuck Goodwin's work has made us exquisitely sensitive to these "one-to-many" possibilities. His work with Chil gives eloquent testimony to the human capacity to mean much with little (Goodwin 2018). Using three simple utterances – "yes, no, and" – Chil can communicate almost anything. Chil pushes the human capacities for ostension and inference to their limits, as he must leverage every available affordance of human communication to make his meaning known. Chuck's analysis of conversations with Chil demonstrates the hard collective work involved in refining that meaning. And we underline that the challenge for Chil and his interlocutors isn't a poverty of meaning – it is an excess of *possible* meaning. This excess highlights strategies and semiotic resources for managing the multiplicity of meaning.

The word "inference" suggests a particularly clean formulation of the problem. At any given moment, parties in an interaction entertain (not always consciously) hypotheses about the possible meanings behind some communicative action. Inference is, then, the use of evidence (or "data") to update the distribution of belief over hypotheses. The challenge facing humans (and any other communicators who possess an ostensive-inferential system) is this: *how to reduce the space of possible hypotheses from vast to manageable?* At this level of abstraction, challenges as diverse as learning words, understanding gestures, and sharing humour converge.

How, then, do humans narrow the hypothesis space of possible meanings? Here we outline five distinct resources that help constrain the meaning space. We do not give an exhaustive treatment to any resource; each has been described in detail by previous scholars, and often supports a vast literature. Rather, we present a brief sketch of the ways in which these resources work together to create meaning in human communication. Wherever possible, we ground the discussion in examples from Chuck's work and connect it to the wide-ranging discussions

[3] It is also the case that multiple signals can map onto the same meaning in the human communication system, but this poses a much less substantial problem than the multiplicity of possible meanings for a *given* signal. While both the ape and the human systems are *technically* many-to-many, the formal essence of the ape system is many-to-one and the human, one-to-many. See https://en.wikipedia.org/wiki/Many-to-many_(data_model)

that characterize his lab group. The communicative resources we propose are: (1) form of the communicative act; (2) physical context of communication; (3) joint activity; (4) shared history; and (5) linguistic convention. Physical context, joint activity, and shared history all fall under what Tomasello (2010) describes as common ground; we find it useful to tease apart these aspects of common ground, as they involve different cognitive mechanisms, may have distinct evolutionary histories, and likely require distinct treatment in computational models of communicating agents.

The **physical form** of the communicative act (hereafter, signal) can reduce the space of possible meaning. The most prominent example is no doubt iconicity, which is prevalent in both gesture (e.g., flapping the hands like a flying bird) and speech (e.g., "sound symbolism" like the famous *bouba-kiki* contrast, Ramachandran, Hubbard 2001; Dingemanse *et al.* 2015). It's easy to see how iconicity helps the recipient; on a connectionist account, for example, certain features of the signal may cause some of the same neurons to fire as would fire in the presence of the referent (see Dingemanse *et al.* 2015 for a discussion of other mechanisms). For the signal producer, however, the analogical challenge appears immediately: she must see how to pick out certain features of the intended referent and produce a signal that shares those features. With the exception of certain cases where this is especially transparent (e.g., using an associated action to refer to an activity by metonymy), exploiting the detailed form of the signal to reduce inferential load is probably beyond the analogical reasoning of early human ancestors, raising doubts for a "pantomimic" theory of language origins (Cartmill, Goldin-Meadow 2012). For an adult human with fully developed ostensive and inferential powers, however, the detailed form of the signal is a rich resource. For example, consider Chil's use of prosody to express his stance toward Chuck's announcement that he has "ordered a hospital bed" for Chil. Chil exploits the iconic (and ironic) relation between his and Chuck's prosody to stake out an unambiguous position (the hospital bed is mentioned in Goodwin 2018: 65; Chuck shared a video of this remarkable encounter with his lab in Summer 2017).

The **physical context** of communication sets up a literal space of possible referents; hence, the details of that space can substantially reduce the space of hypothesized meanings. For example, gestures can exploit this shared space to create meaning by linking to, overlaying, and building off of the physical space, in what Chuck labels environmental coupling (Goodwin 2007). Such coupling between communicative acts and the environment has implications for the comprehension and learning of spoken language. And the coupling need not be gestural: Parents of toddlers whose speech is more transparently coupled to the environment (e.g., by referring to things in the immediately perceivable environment, or by timing the use of a word to coincide with joint attention on an object) have children whose vocabularies are larger by the time they enter school (Cartmill *et al.* 2013). Without necessarily planning to do so, these parents are reducing the space of likely hypotheses for their listeners. Environmental coupling is often supported by the indexicality of a communicative act. The speaker might point to an object in the environment, as we often see in Chuck's data with geologists or oceanographers.

But pointing isn't always transparent, as Chuck demonstrates when analysing Chil's communication about the grapefruit (Goodwin 2018). Indeed, this last analysis underlines the limitation of physical context as a meaning-reducing resource: it works best when the intended meaning is immediately present in a relatively simple environment. Once referents are "out of frame" for one party, indexing features of the physical environment may lead to incorrect inferences about the space of possible meanings. In the grapefruit discussion, for example, Chil's point indexes Candy, whom he knows to be "within the scope of [his] second point" – but Chuck does not, and Chuck's inference about the space of possible meanings is distorted as a result. Chuck believes that Chil wants him to take grapefruit from New Jersey to California, whereas in fact "Chil want[s] Chuck to offer some of the grapefruit to Candy" (Goodwin 2018). In this case, physical context cannot solve Chil's problem – but other resources could, if they were available.

Chuck's work is, perhaps, most distinctive in emphasizing the power of **joint activity** to clarify the space of possible meanings (this occurs frequently in his analysis of scientific work, for example). When individuals are engaged in action together, the space of relevant referents, meanings, and actions contracts (or expands) as an unfolding function of that action. To give a prosaic example: at a recent holiday party, one of us was able to recruit a colleague to move a bottle off a tablecloth by gently tugging on the edge of the tablecloth; the transparency of the ongoing activity ("cleaning up") made the intended meaning obvious. Insofar as the joint activity has a known and legible logic – insofar as it can be projected into the past or the future – joint activity can also clarify meaning in communication about absent, past, future, or imagined matters. We expect that the reduction of meaning through joint activity was an important aspect of the evolution of human language; frequent joint activity (e.g., tool making or food preparation) could provide ample scaffolding (Caporael *et al.* 2013) for relatively simple signal systems to support inference to rich meaning. As joint activities grow more complex, however, multiplying potential actions leads to multiplying potential meanings, and other resources for managing meaning must be brought to bear.

Shared history is one such resource. Humans can draw on the immediate past of an ongoing interaction (available in memory) as a joint resource for meaning-making. This maps onto Chuck's famous "re-use with transformation" argument (Goodwin 2018). For example, one speaker might use gesture to establish a shared "imagined space," within which an interlocutor can refer to or modify entities and events previously established in gesture. Humans also have a remarkable capacity to quickly develop dyadic or small group conventions on the basis of their shared experience – to agree on a particular, idiosyncratic mapping between signal and meaning (like the nicknames, terms of endearment, and private conventional gestures shared by romantic partners). One might think that this capacity is a "free rider" on a more general capacity to learn linguistic conventions; comparison with the pair-specific "bonding rituals" of our distant capuchin monkey relatives, however, suggests that this capacity to create pairwise or small group conventions may be phylogenetically older than and foundational to the learning of linguistic conventions (Perry *et al.* 2003). And two parties with shared history need not

have interacted with each other in the past; for example, the vast majority of our readers already know that Chil is Chuck's father and that he suffered a massive stroke in 1979, because they share with us a history of reading Chuck's work and seeing him talk about and analyse his recordings of Chil. That shared history quickly (and unobtrusively) reduces the space of possible meanings.

Once language is acquired – either phylogenetically or ontogenetically – **linguistic convention** does most of the work of reducing meaning. To first order, we have the (largely arbitrary) relation between a specific word and its possible meanings. To second order, those possible meanings evolve – and meaning-making unfolds – in the larger syntactic context of a linguistic sequence. For a given word, the surrounding speech provides constraints on possible meaning (this is critical during language learning, Gleitman *et al.* 2005). These structural constraints are sufficiently powerful that linguistic convention can support successful communication in the absence of the other four meaning-reducing resources (example: the communication we're hopefully achieving right now in this text without the benefit of iconicity, shared physical space, and joint activity – and with only limited shared history).

The linguistic conventions of human language allow virtually unbounded communication about events, individuals, or objects distant in space or time – even wildly imaginary entities that have never been described before. It is in this sense, then, that the conventional linguistic code *does* expand the expressive power of the O-I system, as Scott-Phillips claims. At the same time, linguistic convention is the most powerful resource for *constraining* that expressive power – or perhaps it's better to say, for *directing* it. To return to the grapefruit example (Goodwin 2018), if Chil had had a full lexicon of conventionalized linguistic symbols at his disposal, he could have easily communicated his intended meaning to Chuck; with the resources at hand, however, it was not possible to reduce the space of possible meanings correctly.

The careful reader could find threads of the arguments we present above woven throughout Chuck's work (above all, in his magnum opus *Co-Operative Action*). These threads are entangled with topics far beyond the study of interaction. For those who study the evolution of language, our emphasis on inferential challenges and the different resources for solving them suggests that the field should re-orient toward those resources especially likely to assist our ancestors (like joint activity and shared history). For those actively working to build artificial intelligences, our typology of resources that contribute to inference in human communicative understanding suggests potentially fruitful lines of investigation: ostension, language-in-interaction, language-in-activity, the formation of small-group conventions, and more. Likewise, the need to "think computationally" about ostension and inference (Valiant 2013; Foster *under review*) motivates the dissection of common ground into cognitively (and computationally) distinct mechanisms.

We close with a final observation. For us, as participants in Chuck's lab, nothing has created greater appreciation for human ingenuity than watching Chuck's videos of Chil. These documents show how many meanings can be made

from just a few words; or rather, how impoverished our view of communication is, if we focus on anything less than the rich interplay of words and worlds, histories and bodies. To understand Chil, Chuck radically expands his sense of what might be meaningful – and of what those meanings might be. Did his interactions with Chil create Chuck's remarkable capacity to let meanings multiply – "to see a World in a Grain of Sand," as Blake put it? We cannot say with any certainty; it may be, instead, that this capacity made Chuck such an effective *interlocutor* for Chil (as suggested by Chuck's moving description in Goodwin 2018: 63). We only know that *our* interactions with Chuck have taught us to see anew; to see more generously and more generatively. We hope to one day make that seeing multiply, as we strive to pass on to our students and interlocutors Chuck's alchemical vision: a vision that makes the world a site of constant wonder and infinite meaning.

References

Byrne, Richard W.; Cartmill, Erica A.; Genty, Emilie; Graham, Kirsty E.; Hobaiter, Catherine; Tanner, Joanne E. 2017. Great ape gestures: Intentional communication with a rich set of innate signals. *Animal Cognition* 20(4): 755–769. https://doi.org/10.1007/s10071-017-1096-4.

Caporael, Linnda R.; Griesemer, James R.; Wimsatt, William (eds.) 2013. *Developing Scaffolds in Evolution, Culture, and Cognition*. Cambridge, Massachusetts: MIT Press.

Cartmill, Erica A. 2016. Lasting impacts of the code model on primate communication research. In: Roberts, Sean; Cuskley, Christine; McCrohon, Luke; Barceló-Coblijn, Lluis; Fehér, Olga; Verhoef, Tessa (eds.), *The Evolution of Language: Proceedings of the 11th International Conference (EVOLANG11)*. Available online: http://evolang.org/neworleans/papers/170.html

Cartmill, Erica A. 2015. Comment on "Nonhuman primate communication, pragmatics, and the origins of language." Current Anthropology 56(1): 66–67. https://doi.org/10.1086/679674.

Cartmill, Erica A.; Byrne, Richard W. 2010. Semantics of primate gestures: Intentional meanings of orangutan gestures. *Animal Cognition* 13(6): 793–804. https://doi.org/10.1007/s10071-010-0328-7.

Cartmill, Erica A.; Goldin-Meadow, Susan 2012. Is pantomime a likely stage in language evolution?: Evidence from human and primate gesture. In: Scott-Phillips, Thom C.; Tamariz, Mónica; Cartmill, Erica A.; Hurford, James R. (eds.), *The Evolution of Language: Proceedings of the 9th International Conference (EVOLANG9), Kyoto, Japan, 13–16 March 2012.* Singapore: World Scientific, 418–419.

Cartmill, Erica A.; Armstrong, Benjamin F.; Gleitman, Lila R.; Goldin-Meadow, Susan; Medina, Tamara N.; Trueswell, John C. 2013. Quality of early parent input predicts child vocabulary 3 years later. *Proceedings of the National Academy of Sciences* 110 (28): 11278–11283. https://doi.org/10.1073/pnas.1309518110.

Dingemanse, Mark; Blasi, Damián E.; Lupyan, Gary; Christiansen, Morten H.; Monaghan, Padraic 2015. Arbitrariness, iconicity, and systematicity in language. *Trends in Cognitive Sciences* 19(10): 603–615. https://doi.org/10.1016/j.tics.2015.07.013.

Foster, Jacob G. under review. Culture and computation: Steps to a probably approximately correct theory of culture.

Gleitman, Lila R.; Cassidy, Kimberly; Nappa, Rebecca; Papafragou, Anna; Trueswell, John C. 2005. Hard words. *Language Learning and Development* 1(1): 23–64. https://doi.org/10.1207/s15473341lld0101_4.

Goodwin, Charles 1994. Professional vision. *American Anthropologist* 96(3): 606–633. https://doi.org/10.1525/aa.1994.96.3.02a00100.

Goodwin, Charles 2007. Environmentally coupled gestures. In: Duncan, Susan D.; Cassell, Justine; Levy, Elena T. (eds.), *Gesture and the Dynamic Dimension of Language*. Amsterdam/Philadelphia: John Benjamins, 195–212.

Goodwin, Charles 2018. *Co-Operative Action*. Cambridge: Cambridge University Press.

Levine, Donald N. 1988. *The Flight from Ambiguity: Essays in Social and Cultural Theory*. Chicago: University of Chicago Press.

Perry, Susan; Baker, Mary; Fedigan, Linda; Gros-Louis, Julie; Jack, Katherine; Mackinnon, Katherine C.; Manson, Joseph H.; Panger, Melissa; Pyle, Kendra; Rose, Lisa 2013. Social conventions in wild white-faced capuchin monkeys: Evidence for traditions in a neotropical primate. *Current Anthropology* 44(2): 241–268. https://doi.org/10.1086/345825.

Quine, Willard Van Orman 1960. *Word and Object*. Cambridge, Massachusetts: MIT Press.

Ramachandran, Vilayanur S.; Hubbard, Edward M. 2001. Synaesthesia: A window into perception, thought and language. *Journal of Consciousness Studies* 8(12): 3–34.

Scott-Phillips, Thom C. 2015. Nonhuman primate communication, pragmatics, and the origins of language. *Current Anthropology* 56(1): 56–80. https://doi.org/10.1086/679674.

Shanahan, Murray 2016. The frame problem. In: Zalta, Edward N. (ed.), *The Stanford Encyclopedia of Philosophy*. California: Stanford University. https://plato.stanford.edu/archives/spr2016/entries/frame-problem/.

Sperber, Dan; Wilson, Deirdre 1996. *Relevance: Communication and Cognition*. New Jersey: Wiley.

Tomasello, Michael 2010. *Origins of Human Communication*. Cambridge, Massachusetts: MIT Press.

Valiant, Leslie 2013. *Probably Approximately Correct: Nature's Algorithms for Learning and Prospering in a Complex World*. New York: Basic Books.