THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Linguistic alignment between people and computers

Citation for published version:

Branigan, HP, Pickering, MJ, Pearson, J & Mclean, JF 2010, 'Linguistic alignment between people and computers', *Journal of Pragmatics*, vol. 42, no. 9, pp. 2355-2368.
<https://doi.org/10.1016/j.pragma.2009.12.012>

Digital Object Identifier (DOI):

[10.1016/j.pragma.2009.12.012](https://doi.org/10.1016/j.pragma.2009.12.012)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Journal of Pragmatics

Publisher Rights Statement:

© Branigan, H. P., Pickering, M. J., Pearson, J., & Mclean, J. F. (2010). Linguistic alignment between people and computers. *Journal of Pragmatics*, 42(9), 2355-2368. 10.1016/j.pragma.2009.12.012

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Contents lists available at ScienceDirect

Journal of Pragmatics

journal homepage: www.elsevier.com/locate/pragma

Linguistic alignment between people and computers

Holly P. Branigan^{*}, Martin J. Pickering, Jamie Pearson, Janet F. McLean

Department of Psychology, University of Edinburgh, 7 George Square, Edinburgh EH8 9JZ, United Kingdom

ARTICLE INFO

Article history:

Received 17 December 2009

Accepted 23 December 2009

Keywords:

Alignment

Human–computer interaction

Dialogue

Beliefs

Syntax

Lexicon

ABSTRACT

There is strong evidence that when two people talk to each other, they tend to converge, or align, on common ways of speaking (e.g., [Pickering and Garrod, 2004](#)). In this paper, we discuss possible mechanisms that might lead to linguistic alignment, contrasting mechanisms that are encapsulated within the language processing system, and so unmediated by beliefs about the interlocutor, with mechanisms that are mediated by beliefs about the interlocutor and that are concerned with considerations of either communicative success or social affect. We consider how these mechanisms might be implicated in human–computer interaction (HCI), and then review recent empirical studies that investigated linguistic alignment in HCI. We argue that there is strong evidence that alignment occurs in HCI, but that it differs in important ways from that found in interactions between humans: It is generally stronger and has a larger mediated component that is concerned with enhancing communicative success.

© 2010 Elsevier B.V. All rights reserved.

A conversation is a joint activity between interlocutors, whose contributions are tied in with their partners' contributions ([Clark, 1996](#)). An isolated speaker¹ producing a monologue has to work on her own model of what to say, together with assumptions about the presumed audience. But an interlocutor's contribution in a dialogue is dependent on prior contributions by other interlocutors, for instance in response to a question or a request for more information. Importantly, this dependency relates not only to the content, but also to the form of contributions ([Linell, 1998](#)). In this paper we discuss the mechanisms that might underlie such *alignment* in interactions between two people, and then consider how these mechanisms might affect the way in which people interact with computers.²

1. Evidence for alignment in dialogue between people

Dialogue involves the convergence of behavior between interlocutors at many different levels. At non-linguistic levels, interlocutors imitate each other in many respects, including facial expression ([Bavelas et al., 1986](#)) and other indicators of their emotional states (e.g., [Hatfield et al., 1994](#)). But qualitative analyses of repetition in naturalistic corpora ([Tannen, 1987](#)), field experiments ([Giles and Powesland, 1975](#)), extensive quantitative studies (e.g., [Gries, 2005](#)) and laboratory studies demonstrate that such convergence is also prevalent at linguistic levels.

In this paper, we say that interlocutors who converge on linguistic behavior have *aligned* their behavior, and use the term *alignment* to refer to such convergence. Thus, two interlocutors who both use *sofa* to refer to a type of seat are aligned on the

^{*} Corresponding author. Tel.: +44 131 650 8387; fax: +44 131 650 3461.

E-mail address: holly.branigan@ed.ac.uk (H.P. Branigan).

¹ We will use the term *speaker* to refer to any agent that produces language (whether speech or text), and will adopt the convention of assuming a female speaker and a male addressee.

² We will use the term *computer* to refer to any artificial system.

use of that term, whereas two interlocutors who use different terms, such as *sofa* and *couch*, are not aligned. However, *alignment* strictly refers to interlocutors' mental representations, not their actual behavior (Pickering and Garrod, 2004; see Costa et al., 2008). For example, two interlocutors can be aligned on the use of *sofa* if they would tend to call such a seat a *sofa*, even if one or both of them never actually utters the term.

Some such alignment relates to the form of interlocutors' contributions. For example, interlocutors tend to align on accent and speech rate (e.g., Giles et al., 1991), and adopt increasingly similar phonetic realizations of repeated words (Pardo, 2006). Interlocutors also tend to repeat each other in terms of their grammatical choices. Levelt and Kelter (1982) telephoned (Dutch) shopkeepers and found that they tended to say *At 5 o'clock* in response to *At what time does your shop close?*, but tended to say *5 o'clock* in response to *What time does your shop close?*. Branigan et al. (2000) used a confederate-scripting paradigm in which participants took turns to describe and match picture cards; one participant was a confederate who produced descriptions following a script. Branigan et al. found that the naïve participant tended to use the form of utterance just used by the confederate even when no content words were repeated. For example, they tended to use a "prepositional object" (PO) form like *the pirate giving the book to the swimmer* following another prepositional object sentence but a "double object" (DO) form like *the pirate giving the swimmer the book* following another double object sentence. Similar alignment also occurs with other denotationally equivalent but syntactically distinct constructions (Cleland and Pickering, 2003; Haywood et al., 2005), and even between languages in bilinguals (Hartsuiker et al., 2004). Overall, there is very strong evidence that interlocutors align their syntactic representations (see Pickering and Ferreira, 2008).

Alignment also occurs at levels of structure related to meaning, in which the choice of one alternative over another is associated with different semantic representations. For example, speakers may align on particular ways of describing spatial positions. Garrod and Anderson (1987) had pairs of participants play a cooperative maze game, in which they took turns to describe their positions to each other. There are many different ways to describe one's position, but participants tended to align on the same description scheme. For example, if one player said *I'm two along, four up*, her partner tended to say *I'm one along, five up*; whereas if she said *I'm at B4*, her partner tended to say *I'm at A5*. These players aligned on a "path" or a "coordinate" description scheme, rather than specific words. They also aligned on the interpretation of these descriptions, for example treating the origin as the bottom left corner of the maze. Other evidence shows alignment of other aspects of spatial descriptions, such as the interpretation of left and right (Schober, 1993; Watson et al., 2004).

Perhaps the most extensive evidence for alignment relates to word choice (sometimes called *lexical entrainment*). For example, Garrod and Anderson (1987) found that interlocutors used the same expressions, and moreover that they tended to interpret words in the same way (e.g., *line* to mean horizontal row of nodes in the maze). Brennan and Clark (1996) had directors describe a set of cards depicting common objects to matchers so that they could reconstruct the directors' array. One set of trials contained multiple objects from the same category. Directors and matchers settled on subordinate-level terms to refer to the objects (e.g., *pennyloafer*), because basic-level terms (e.g., *shoe*) would not discriminate between these objects. A subsequent set of trials included one object from each category, so basic-level terms would now be sufficient. However, participants often continued to use the subordinate terms. Moreover, such effects are partner-specific, so that speakers will consistently use the same term to a particular interlocutor, and a particular interlocutor will expect the speaker to use the same term to him, such that he experiences difficulty if the speaker does not consistently use that term (Metzing and Brennan, 2003). Horton and Gerrig (2005) suggested that such partner-specific effects might fall out of normal memory processes, in which interlocutors represent as a conjoint cue information about linguistic expressions and the person with whom the expression was used, so that the presence of that particular interlocutor activates that expression.

2. The role of alignment in successful communication

Clearly convergence of both non-linguistic and linguistic behavior is robust and pervasive in dialogue. According to Pickering and Garrod (2004), such convergence is in fact instrumental in bringing about successful communication. They argued that interlocutors understand each other when they align their model of the situation under discussion, and that such alignment is largely the result of alignment at different levels of linguistic representation. Alignment of some levels of structure clearly corresponds to alignment in ways of perceiving the world. For example, by converging on a path description scheme, interlocutors in Garrod and Anderson's (1987) experiment were also converging on a common way of semantically representing the maze. Similarly, people who use the same expression to refer to an object or concept are likely to take on board the same connotations associated with the expression, as is apparent from occasions when adversaries refuse to align their expressions (e.g., advocates' use of *fetus* vs. *unborn child*; Danet, 1980). It is not surprising, then, that alignment at these levels is critical to convergence on common situation models. But Pickering and Garrod suggested that alignment of representations at levels that are not meaning-based (e.g., syntax, pronunciation) also plays a fundamental role in alignment of semantic representations. In their account, alignment of each level of linguistic representation enhances alignment at other levels, so that for example syntactic alignment enhances semantic alignment. Alignment of linguistic representations at all levels therefore plays a causal role in the attainment of mutual understanding: Failure to align at different levels of language may result in failure to communicate successfully.

Other researchers have questioned the causal link between alignment and successful communication, specifically the proposal that alignment at one level of structure leads to alignment at other levels. For example, Garrod and Clark's (1993) findings of failures in communication when children aligned at a lexical level but failed to align at a semantic level suggest

that lexical alignment does not always result in semantic alignment, and more generally that alignment need not lead to communicative success.

3. What mechanisms underlie alignment?

Although it is uncontroversial that behaviors leading to alignment are pervasive in interactions between humans, more than one mechanism has been proposed to account for these behaviors. We will focus below on mechanisms that have been proposed with respect to linguistic alignment.

3.1. Alignment as an unmediated mechanism

One approach to alignment appeals to the priming of representations and processes that are not affected by extralinguistic information. Extensive research has shown that linguistic processing at many levels is strongly influenced by prior processing of related material. For example, interpreting *pennyloafer* requires activating the *pennyloafer* lexical entry (e.g., Meyer and Schvaneveldt, 1971); processing a sentence using the PO construction requires activating the grammatical rule or rules associated with that construction (Branigan et al., 2000). Standard accounts of language processing assume that activation of such representations or processes does not decay immediately, and the persistent activation increases the likelihood that the comprehender will subsequently use the associated linguistic form (e.g., Pickering and Branigan, 1998). The activation from comprehension and the effects of activation on subsequent behavior are necessary consequences of the architecture of the language processor. Such priming is not mediated by the speaker's beliefs about her interlocutor. We will therefore term these mechanisms *unmediated* mechanisms.

An unmediated component to alignment would be consistent with the evidence that interlocutors may be almost entirely unaware that alignment has taken place (Pickering and Garrod, 2004). It almost always arises without explicit negotiation, and on those occasions where speakers do explicitly negotiate a term to use, they frequently end up aligning on a different expression (Garrod and Anderson, 1987). Furthermore, speakers are usually unaware of aligning with a conversational partner. Post-experimental debriefing has shown that speakers very rarely notice alignment of form; they sometimes, though not frequently, report awareness of alignment at levels related to meaning. At phonological and acoustic levels alignment can occur very rapidly and seems largely resource-free (Fowler et al., 2003), with listeners activating appropriate muscles in the tongue while listening to speech but not during non-speech (Fadiga et al., 2002; see also Watkins and Paus, 2004). Additionally, speakers that shadow spoken words produce more faithful imitations of the original when they speak more rapidly (Goldinger, 1998). Evidence from language acquisition is also consistent with a strong unmediated component to alignment: Children show a stronger tendency to align than adults under at least some circumstances; notably, they align linguistic form even when this leads to misunderstanding, such as using the same term with different reference (e.g., using *square* to mean different things; Garrod and Clark, 1993). Garrod and Clark used such evidence to propose that children align as their default behavior, and that part of becoming a mature language user involves learning to suppress the tendency towards alignment when necessary.

3.2. Alignment as audience design

Pickering and Garrod (2004) suggested that alignment in dialogue is largely based upon unmediated mechanisms, primarily because of the cognitive economy that they offer, since they do not require the speaker to model her interlocutor. But there are also good reasons why we might expect repetition to be mediated by a speaker's beliefs, so that a speaker might choose to use a particular expression because she believes that it is the appropriate expression to use for that particular interlocutor. Most obviously, speakers might choose to use expressions that they believe will enhance communicative success, in a process of *audience design* (Bell, 1984). For example, if a speaker has a choice of two terms, she may choose whichever term she believes her interlocutor is most likely to understand.

Clark (1996) suggested that speakers make use of two types of evidence when assessing what their interlocutor is likely to know, and hence to understand. The first is evidence about the *cultural communities* (i.e., cultural groups, systems or networks), to which the speaker believes that the addressee belongs. For example, if the speaker believes that her interlocutor is a fellow native of Edinburgh, she may assume that the interlocutor has access to certain specialized knowledge about local geography, and so may use a proper name to refer to a particular landmark (e.g., *The Scott Monument*), whereas if the speaker believes that her interlocutor is a visitor from a different city, she may tend to use a longer and more detailed description (*The tall pointy monument that looks like Thunderbird 2*). Experimental studies show that speakers accordingly tailor the content of their utterances to match their *a priori* assumptions about the social distribution of knowledge, for example the fact that students are more likely to be familiar with film stars than with industrialists (Fussell and Krauss, 1992).

Moreover, such beliefs can affect the form of speakers' utterances, as well as their content. For example, beliefs about the linguistic competence of the addressee may cause the speaker to speak more slowly or to use less complex syntax to a non-native speaker than when addressing another native speaker (Ferguson, 1975). Such effects underline the fact that speakers' beliefs about their interlocutors' community membership include beliefs about the speech community to which they belong, and so their linguistic knowledge: what sorts of linguistic representations are available to them. In these cases, the speaker

draws upon her beliefs about what a member of the community to which she believes her addressee belongs is likely to know.

The second type of evidence to which the speaker may make reference is her own direct personal experience with that interlocutor. The speaker may assume that she shares certain perceptual experiences with her interlocutor, or may refer to the language that she and her interlocutor have already used. If she has used the term *chair* to refer to an object, and her interlocutor appears to have understood it, then she may assume that this is a good term to use: *I believe you understand expression X because you have already given evidence that you understand expression X*. Similarly, if her interlocutor has used to the word *chair* to refer to the same object, she may again assume that this is a good term to use: *I believe you understand expression X because I have already seen you produce expression X*. Both beliefs would encourage her to re-use the expression. This may explain Bortfeld and Brennan's (1997) finding that native speakers of English tended to align with non-native speakers' use of highly unidiomatic expressions (e.g., *The chair that can go back and forth* to refer to a rocking chair) when the non-native speaker gave evidence of difficulties in comprehending.

Direct personal evidence, through the content or form of the interlocutors' contributions, can also dynamically affect speakers' beliefs about their interlocutors' community membership (and hence presumed knowledge), such as an interlocutor's comment that he is from Manchester, his accent, or his apparent failure to understand a particular term (e.g., Isaacs and Clark, 1987). Throughout a dialogue, then, speakers may continually be adjusting their beliefs about their interlocutors, and this adjustment may be very rapid. In the case of Bortfeld and Brennan's (1997) participants, a non-native speaker's failure to understand the term *rocking chair* might lead a native speaker to adjust her beliefs about the sorts of linguistic expressions that he might understand, as well as causing her to align with his choice of term.

Mechanisms that make reference to beliefs about the interlocutor would serve to enhance the likelihood that the content of the speaker's message would be correctly understood, and would therefore increase the likelihood of communication being successful. Such mechanisms are plausibly implicated in alignment of aspects of language associated with differences in meaning (e.g., lexical choice). But we might expect such a mediated component to be relatively less important in normal dialogue between two adult native speakers than the unmediated priming-based component, because there are no obvious reasons why communication should be unsuccessful, and because such modelling of the interlocutor is cognitively costly (e.g., Horton and Keysar, 1996).

3.3. Alignment for social affect

The repeated use of expressions might also serve a different strategic purpose, related to the social relationships between the participants in the interaction. Speakers might repeat their interlocutors' choice of expressions in order to increase the interlocutor's positive feelings towards the speaker. Speakers might also repeat expressions as an expression of affiliation with the interlocutor: Because they have positive feelings towards their interlocutors, they want to behave like them (irrespective of whether this will cause the interlocutors to feel positively towards them in return). Certainly there is now a large body of evidence that the ways in which speakers align with each other is influenced by social factors such as power relationships, community membership and so on (see Giles et al., 1991, for discussion). More generally, alignment may be subserved by social norms relating to interaction, such as the principle of reciprocity (Gouldner, 1960).

Critically, many studies have demonstrated that repeating an interlocutor's behavior in dialogue has affective consequences. With respect to non-verbal behavior, Chartrand and Bargh (1999) demonstrated that confederates were rated more favorably by their interlocutors when they mimicked the posture and movements of the interlocutor than when they did not, even when the interlocutor was not aware of the mimicry (see also Lakin and Chartrand, 2003). Similarly, Maurer and Tindall (1983) found that counselors who mimicked the body postures of their clients were rated as more empathetic. Comparable effects are found with respect to alignment of language. For example, speakers who converge with respect to breadth of vocabulary (Bradac et al., 1988) are judged more favorably than those who do not. Such positive affect can have very tangible consequences: van Baaren et al. (2003) found that waitresses received larger tips when they repeated their customers' orders exactly than when they paraphrased them. Furthermore, the positive affect extends beyond the specific interaction. People who have been mimicked display greater prosocial behavior (e.g., helpfulness and generosity) towards other people who were not involved in the mimicry interaction (van Baaren et al., 2004). van Baaren et al. (2004) therefore suggested that imitating each other's behavior may therefore have an evolutionary basis, in that mimicry has the function of acting as 'social glue' that helps to strengthen social bonds (see also Lakin et al., 2003).

We might therefore expect that speakers' choice of expressions reflects social factors. If an interlocutor has already used the term *chair*, for example, the speaker may be conventionally expected to use the same term; using an alternative term would be impolite, as it would imply that the interlocutor's choice of term was in some way inappropriate or wrong, and might result in the interlocutor experiencing negative affect towards the speaker. The speaker would therefore use an expression if the interlocutor has already used that expression because it is polite and likely to cause positive affect.

3.4. Alignment from multiple components

Although most investigations of alignment and imitation effects tend to adopt one or other of these approaches to explain their findings, these different types of explanation are of course not mutually exclusive: Alignment in a particular context may have both unmediated and mediated components. For example, unmediated alignment might occur alongside

mediated alignment that is designed to enhance communicative effectiveness, just as unmediated alignment may occur alongside audience design concerned with the avoidance of syntactic ambiguity (Haywood et al., 2005). The balance between these distinct components might differ between communicative contexts. For example, strategic components may play a greater role in contexts where effective communication is particularly important or there are reasons to believe that it may not be successful, for example because one participant has a lower level of linguistic competence than the other (see Costa et al., 2008). Equally, the different components might affect alignment at some levels of structure more than at others. For example, in the same way that levels of structure associated with differences in meaning appear to be more amenable to audience design effects, they might also be more amenable to strategic alignment effects that are mediated by beliefs.

4. Why study alignment in human–computer interaction?

So far we have considered the evidence for alignment in interactions between humans, and the possible components that might underlie these effects. We now consider whether alignment might also occur in interactions between humans and computers, and the way in which any such effects might manifest themselves. These are important issues for two reasons. First, the increasing place of technology, and in particular the emergence and rapid development of natural language interfaces to dialogue systems, mean that dialogues between people and artificial systems are becoming a part of everyday life (e.g., Nass and Moon, 2000). Investigating the characteristics of such dialogues will help our understanding of how human–computer interaction works and will aid the development of robust naturalistic systems. It may also elucidate potential limits on human–computer interaction: If alignment is critical for effective communication, as Pickering and Garrod (2004) argued, but alignment does not occur (or only occurs to a restricted extent) in HCI, then this would imply that HCI is inherently a less effective context for communication than dialogue between humans.

Additionally, studying alignment in HCI may cast light on alignment phenomena in general, and hence inform our models of dialogue. In particular, as we shall see below, studies of alignment in HCI can clarify the extent to which alignment is mediated in some way by beliefs. In normal dialogue, it is very difficult to control all of the various parameters that might be relevant to alignment. But HCI makes it possible to manipulate beliefs independently of other aspects of the interaction, so that a situation can be established in which a speaker believes she is carrying out a task by interacting with another person or with an artificial system that is able to perform the task but is in other respects clearly not human, and in which the behavior that she encounters under either scenario is identical. By examining the extent to which speakers align under different belief conditions when the interlocutors' actual behavior is held constant, it is possible to isolate the contribution that beliefs make to alignment, and moreover to investigate the factors that may underlie the establishment and revision of these beliefs.

5. How might alignment occur in HCI?

We can make very different predictions about alignment effects in interactions between humans and computers, depending on the mechanisms that we assume to be operative. If only the unmediated component is relevant, then alignment should be determined by the behavior of the interlocutor and should be unaffected by the speaker's beliefs about that interlocutor. In contrast, if alignment is underpinned by mediated mechanisms, then speakers should align to a greater or lesser extent depending on the extent to which they believe their interlocutor to be able to benefit communicatively or to be an agent that can experience social affect. As for dialogues between humans, it is possible that there are both unmediated and mediated components of alignment in HCI.

We can identify two aspects to alignment in HCI, reflecting the two directions of communication: the computer's contributions directed towards the user, and the user's contributions directed towards the computer. The first aspect therefore relates to computers that align with a human user's contributions, and how such computers affect the user. The second aspect relates to humans' alignment with a computer's contributions.

5.1. Effects of alignment by computers

Unmediated mechanisms of alignment would affect speakers' linguistic behavior when interacting with a computer that aligns with their previous utterances. This is because such alignment by a computer would cause a 'spiraling' effect, in which the computer's use of an expression would reinforce the activation of that expression in the speaker, leading her to use that expression again in her turn (Branigan et al., 2005). For example, when a speaker uses the term *couch*, it activates the relevant mental representation. If a computer subsequently aligns with the speaker and so uses that term in its own contribution, the speaker's mental representation, here the lexical entry for *couch*, will become even more highly activated by the act of comprehending it. But there is no reason to believe that unmediated mechanisms of alignment would be implicated in how people feel about systems that align with them. It should be irrelevant to the speaker whether or not the system has previously aligned with her.

In contrast, mediated mechanisms that make reference to beliefs about social factors would lead the speaker to behave differently towards computers that align with her than to systems that do not align. In terms of their linguistic behavior, speakers might be more likely to align with a computer that has aligned with them than with one that has not, because of reciprocity effects (as we discuss in section 5.2.1). Speakers should also feel more positive affect when interacting with a computer that aligns with them than with one that does not. Note that such emotions are determined purely with respect to

the user (because it is the user who experiences positive affect); they do not depend upon the characteristics of the computer, specifically whether it is perceived as a social agent that can itself experience positive affect.

5.2. Alignment with computers

5.2.1. Does alignment with computers occur?

A potentially more interesting question concerns humans' alignment with computers' contributions. Here we can make varying predictions with respect to both the existence and the magnitude of alignment, depending on the component processes that are involved. Most fundamentally, the different components differ with respect to whether they would cause people to show alignment with a computer. Mechanisms that lead to alignment through automatic priming of representations and processes might be influenced by considerations internal to the language processing system, for example baseline preferences for the relevant expression. This would be consistent with evidence that highly dispreferred syntactic structures may not be primed (Pickering et al., 2002). In that case we might expect that speakers would not align if the computer produced highly dispreferred expressions. However, they should not be affected by beliefs about the interlocutor, and hence alignment should occur whenever a speaker is exposed to (and attends to) a particular linguistic expression, irrespective of whether that expression appears to come from a person or a computer.

A component of alignment that involves the mediation of beliefs would cause alignment in these circumstances only if the speaker holds particular beliefs about her interlocutor. If the relevant mechanisms are associated with a strategy of communicative enhancement, speakers should align with a computer if they believe that the computer can benefit communicatively from it; for example, they may align with its choice of name for an object if they believe that it can understand that term but not another. This component of alignment would cause speakers to align with a computer's use of a dispreferred or infelicitous expression – overcoming their normal linguistic preferences – if they believed that this was necessary for successful communication, in the same way that a speaker may align with another person in the use of a strongly dispreferred or even inaccurate term in order to communicate successfully (e.g., aligning with a non-native speaker's use of *wheel* to mean tire; Bortfeld and Brennan, 1997). Hence with respect to this component, the important question is whether people believe that computers can benefit from such adaptations. Certainly speakers do adapt their utterances when addressing a computer in ways that are most straightforwardly construed as an attempt to produce more comprehensible utterances; for example, they produce hyperarticulated, slower and louder speech (Oviatt et al., 1998). Such adaptation appears similar to native speaker adaptation to non-native speaker addressees (Ferguson, 1975).

A component that is concerned with social affect, on the other hand, makes reference to beliefs about computers as social agents: Speakers might align with a computer's use of language to the extent that they believe it to be a social agent that is capable of experiencing affect. For example, a speaker might align with a computer's name for an object because she believes that it would be impolite in some sense to use an alternative term. Alternatively, speakers might align with a computer's use of language because they feel an affiliation with the computer and hence wish to emulate it as an expression of their affiliation.

It seems surprising to suggest that speakers might experience such beliefs with respect to such an obviously inanimate and non-sentient entity as a computer, particularly in situations where the non-humanness of the computer is emphasised. Yet there is increasing evidence to suggest that people might hold these beliefs with respect to computers in aspects of their behavior other than alignment. Substantial research has shown that people treat computers as social agents, and as such automatically respond to social cues from them (e.g., Reeves and Nass, 1996); for example, they attribute social characteristics such as personality towards computers (Nass et al., 1995). Moreover, they display behavior towards them as if they were social agents. Nass et al. (1999) showed that people conform to politeness norms when interacting with computers, for example in avoiding direct criticism. Similarly, Fogg and Nass (1997) found that people display reciprocity effects towards computers, performing more and better work for computers that had helped them on a previous task. They also show affiliation effects, rating more highly and as more similar to themselves a computer with which they had to work interdependently to perform a task than a computer with which they did not have to work (Nass et al., 1996). All of these effects would lead us to expect alignment in HCI if alignment arises from considerations of social affect.

5.2.2. How strong is alignment with computers?

The different components of alignment vary not only in whether they would lead to the existence of any alignment in HCI; they also vary in the magnitude of any such effects that they might cause. An unmediated component causes alignment through the activation of linguistic representations and/or processes. Variations in alignment caused by the unmediated component would be due entirely to variations in the encapsulated linguistic system of the speaker. Assuming that people attend to a computer interlocutor to the same extent as they attend to a human interlocutor, we might expect comparable alignment in HCI as in normal human dialogue. Such alignment should occur irrespective of the behavior of the interlocutor with respect to signalling understanding or failure to understand.

A mediated component to alignment would also lead to variations in alignment, but these would be tied to differences in the beliefs that the speaker held about their computer interlocutor. For example, a component associated with communicative enhancement would cause speakers to align more strongly with a computer that they believed would benefit more from alignment than with one that they believed would not benefit so strongly. The critical factors in determining the magnitude of alignment would therefore be those factors that underlie beliefs about the determinants of

communicative success, such as factors governing beliefs about the general or specifically linguistic capability of their interlocutor. In this respect, we might expect more alignment with a computer interlocutor than with a human interlocutor, assuming that humans are generally believed to be more capable than computers in the domain of communication. If those beliefs are dynamically updated in the light of feedback from the interlocutor, as in human dialogue (e.g., Isaacs and Clark, 1987), then speakers might align with a computer to a comparable extent as with a human, if the computer gives feedback of understanding.

Conversely, a component that was mediated by considerations of social affect would cause alignment to vary according to the speaker's desire to bring about positive social affect or to affiliate. In that case, the factors that influence the magnitude of alignment would be those factors that influence the desirability of social affect or affiliation, for example the extent to which the speaker perceives the interlocutor as being of greater social importance. Depending on beliefs, we might expect comparable alignment in HCI and human dialogue (if computers are seen as social agents of comparable importance to humans), or less alignment in HCI than in human dialogue (if computers are seen as inferior social agents compared to humans). There is however no reason to expect that computers should be seen as superior social agents to humans, and hence no reason to expect stronger alignment to computers under this account.

Just as in dialogue between humans, more than one component of alignment might be operative in the same interaction, and furthermore different levels of alignment might be susceptible to different influences. Hence it is possible that HCI involves both unmediated and mediated elements, and that the relative strength of these components in a particular interaction might depend upon the particular characteristics of the communicative context (e.g., features of the computer) and the level of linguistic structure under investigation. In particular, we might expect goal-directed (mediated) effects to manifest themselves more strongly in contexts where successful communication is at a premium or is in some way endangered, and at levels of structure related to meaning (which are more critical to successful communication and more open to conscious decision-making); in contrast, unmediated effects might manifest themselves more strongly when successful communication is not critical or is not under threat, and at levels of structure less related to meaning.

6. Empirical evidence of alignment in HCI

6.1. Evidence of effects of alignment by computers

There has been a recent surge of interest in alignment effects in HCI. Some studies have been concerned with the effects on communication when computers align with their human user. Together, they provide evidence that alignment by a computer induces social affect in a human interlocutor. With respect to non-linguistic alignment, Bailenson and Yee (2005) found that participants rated more positively an animated computer agent that mimicked their head movements than one that did not, much as humans who mimic the body posture of their interlocutors are rated more highly by them (e.g., Chartrand and Bargh, 1999). Moreover, people are sensitive to the dynamic process of alignment: Moon and Nass (1996) demonstrated that participants rated more highly a computer system when its manifested personality altered during the interaction to be more aligned to their own than when it manifested the same (aligned) personality throughout the interaction.

More relevantly, comparable effects have been found for linguistic alignment. Nass and Lee (2001) investigated the effects of alignment of personality, as indexed by four acoustic features (intensity, fundamental frequency, frequency range and speech rate) in a speech-based dialogue system that presented participants with book reviews. Nass and Lee manipulated whether participants interacted with a computer whose synthesized voice manifested a personality that was aligned with their own personality in terms of introversion vs. extroversion. For example, the 'extrovert' had a higher volume, fundamental frequency, frequency range, and speech rate—all characteristics associated with judgments of extroversion. Nass and Lee found that participants liked voices that matched their own personality more than voices that did not. Similarly, Ward and Nakagawa (2002) found that speakers rated a telephony system that aligned with their speech rates more favourably than a system that did not. Just as alignment by a human interlocutor causes people to rate that interlocutor more positively, then, so alignment by a computer induces positive affect towards that computer. These results accord with the evidence that people treat computers as social agents and respond appropriately to their social cues (Nass and Moon, 2000).

6.2. Evidence of alignment with computers

Other research has been concerned with the second aspect of alignment in HCI, namely the question of whether humans spontaneously align with computers.

6.2.1. Do speakers align with computers?

Much of this research has examined alignment with respect to the acoustic-prosodic properties of utterances. The research has yielded consistent evidence that speakers align with a computer interlocutor at these levels of structure. Bell et al. (2003) investigated whether speakers adapted their speech rate when they interacted with an animated character in a simulated spoken dialogue system; in fact, they used a 'Wizard of Oz' paradigm, in which a human experimenter pretended to be the 'computer'. They manipulated whether the system produced feedback with a fast or slow speech rate. They found

that speakers adapted their speech rate to align with that of the computer, producing slow speech in response to the ‘slow computer’ and fast speech in response to the ‘fast computer’. Interestingly, this tendency varied according to the system’s overall behavior: Speakers speeded up less with a ‘fast computer’ when it appeared to be having difficulty comprehending their contributions, and slowed down less with a ‘slow computer’ when it appeared to be handling their contributions successfully. So speakers were sensitive to the feedback that they received, and appeared to adjust the magnitude of their alignment accordingly.

In a series of experiments, Oviatt et al. (2004) examined whether children’s speech aligns with that of a computer interlocutor. Children interacted with a series of animated personae, whose synthesized voices varied in their acoustic properties. They found that children consistently aligned with their animated interlocutor along several different dimensions, including amplitude and pause structure. Such adaptations occurred bidirectionally and dynamically: Children rapidly adapted their speech to align with that of their current interlocutor, for example inserting more pauses when interacting with an ‘introverted’ interlocutor but then inserting fewer pauses when subsequently interacting with an ‘extroverted’ interlocutor, even within the same conversation. As in human dialogue, then, speakers may adapt rapidly to the behavior of a computer interlocutor. More recent work found that adult speakers showed prosodic alignment with a computer interlocutor (Suzuki and Katagiri, 2007).

We can therefore conclude that alignment of at least some aspects of language occurs in HCI, and so that a robust and pervasive feature of human dialogues also characterizes HCI to at least some extent. These findings are reassuring in the respect that if alignment is critical to communicative success, HCI is not thereby doomed to communicative failure. But these findings relate to one particular aspect of language. Does alignment also occur in HCI for other levels of structure, and in particular for those more closely related to meaning (and so which might be more important for successful communication)? Studying such levels of structure might also cast light on the relative contribution of the different potential components of alignment. Although findings of alignment to acoustic-prosodic properties of a computer show that alignment is not dependent on social considerations that apply only to human beings, they do not elucidate whether alignment in HCI is primarily or solely associated with unmediated linguistic mechanisms.

Brennan (1996) reported two experiments that examined lexical alignment in HCI using a Wizard of Oz paradigm. Participants carried out a database query task; on some turns, the ‘computer’ used a different term to refer to the entity that the participant had just referred to. Half of the time, the ‘computer’ explicitly queried the participant’s choice of term and offered the other term as an alternative; the rest of the time, it merely responded using the other term. The participant subsequently had to refer to the same entity, either immediately afterwards or after a delay. Brennan found that speakers tended to adopt the ‘computer’s’ lexical choice after the computer had queried their own choice. But strikingly, they also aligned with the computer lexically 25–80% of the time (depending on condition) when the computer had not queried their term. In both cases, alignment persisted across several intervening references to other objects. These results suggest that alignment in HCI occurs for meaning-related levels of structure. They demonstrate that such effects are persistent and moreover that they can override the normal convention that an initial expression that has been understood and acknowledged to be understood by an addressee will continue to be used by both interlocutors for the remainder of the interaction (Brennan and Clark, 1996; Metzing and Brennan, 2003).

6.2.2. To what extent do speakers align with computers?

In all of the studies above, the key issue of interest was *whether* alignment occurs in HCI. More recent studies have explored *how much* alignment occurs in HCI, and has compared the magnitudes of alignment in HCI versus normal human dialogue. Research of this type offers a more powerful way of addressing the question of which component processes are most relevant to alignment in HCI, and indeed to alignment in dialogue as a whole. One difficulty with comparing alignment in HCI with alignment in interactions between humans is that there are a number of factors that might influence participants’ behavior. For example, the very presence of a computer in the communicative context might be relevant to people’s behavior. A convincing way of solving this problem is to use a paradigm in which all aspects of the interaction are kept the same across conditions except for participants’ beliefs about the nature of their interlocutor. This can be achieved either by using a Wizard of Oz paradigm in which the interlocutor is always a person who pretends to be a computer or a person, depending on condition; or by using what we will term a *reverse Wizard of Oz* paradigm, where the interlocutor is always a computer that is presented to the participants as a computer or a person. In this paradigm, it is possible to control exactly the behavior (linguistic and otherwise) of the interlocutor so that participants always experience the same behavior from their interlocutor; hence any differences in the participants’ behavior when interacting with the interlocutor can only be attributed to differences in their beliefs about their interlocutor.

Using a Wizard of Oz paradigm, Brennan (1991) had participants use a keyboard to ask for information from an interlocutor. Participants were told either that they were interacting via a keyboard with another person (which we will refer to as *computer-mediated interaction*, or CMI) or with a computer (HCI). In fact, their interlocutor was always a human who produced scripted responses. Some aspects of participants’ contributions differed between conditions. For example, speakers produced fewer personal pronouns and fewer acknowledgements when they believed they were interacting with a computer. However, they did not differ with respect to alignment at the level of contribution length: In both the CMI and HCI conditions, people produced short questions when their interlocutor produced short answers, and long questions when their interlocutor produced long answers. This alignment occurred over the duration of the dialogue. At the beginning of the

dialogue, speakers produced longer contributions in the CMI condition and shorter contributions in the HCI condition. But by the second half of the dialogue, there was no difference in contribution length across interlocutor conditions; contribution length was affected only by the interlocutor's response style. Brennan suggested that this showed effects of expectations about the interlocutor (in this case, she suggested, what kinds of linguistic structures the computer could process) and politeness considerations, combined with alignment behavior. In her study, speakers' choice of style was more strongly affected by alignment than by their initial model of the interlocutor.

A recent series of studies has similarly investigated syntactic and lexical alignment in HCI with alignment in CMI conditions using the reverse Wizard of Oz paradigm (Branigan et al., 2003, 2004; Pearson et al., 2006a). All the studies used a modified version of Branigan et al.'s (2000) confederate-scripting paradigm. Participants were led to believe that they were playing a picture-matching and -describing game with an interlocutor in another room via a networked computer terminal, interacting with their unseen interlocutor by typing. On alternate turns, participants described visually presented pictures so that their interlocutor could choose a matching picture, or received descriptions from their interlocutor that matched one of two pictures presented on-screen. Participants believed that they were interacting with either a computer or a human interlocutor, but in fact always interacted with a computer program that produced scripted utterances.

Branigan et al. (2003) investigated how a priori beliefs about an interlocutor affected syntactic alignment. Participants typed descriptions of depicted events for their interlocutor, and in return received descriptions allegedly from their interlocutor to match to pictures presented on-screen. Some of these described ditransitive events. Branigan et al. examined how the syntactic structure of the interlocutor's description (PO vs. DO) affected the syntactic structure that participants produced immediately subsequently to describe an unrelated event involving three entities. They found that participants tended to align with the structure that their interlocutor had just produced, as in the following example of an interaction between the interlocutor and a participant on an experimental trial:

Interlocutor: <i>the chef handing the swimmer the jug</i>	[Prime description, DO structure]
Participant: <i>the pirate showing the sailor the cake</i>	[Target description, DO structure]

This tendency to align occurred whether the verb in the interlocutor's descriptions and the verb in the participant's subsequent descriptions were the same or different, but it was stronger when the verb was repeated than when it was not, as in earlier studies of human dialogue (Branigan et al., 2000). However, alignment was affected by the perceived identity of the interlocutor. When the interlocutor's description and the participant's description involved different verbs, alignment occurred to the same extent for 'human' and 'computer' interlocutors. Hence, participants aligned linguistically with what they believed to be a computer, and the strength of this alignment was broadly comparable with the alignment that occurred when participants believed themselves to be communicating with another person. By contrast, when the interlocutor's description and the participant's description involved the same verb, there was greater alignment to a 'computer' than to a 'human' interlocutor.

These results suggest that syntactic alignment processes in typed dialogue involving no other visible interlocutors are broadly similar to alignment processes in dialogue between co-present interlocutors who use speech to communicate. More interestingly, however, the results help to elucidate the relative contributions of unmediated and mediated mechanisms of alignment in these contexts. The finding of comparable alignment to both 'computer' and 'human' interlocutors when the verb was not repeated can be explained in terms of an unmediated component to alignment, in which alignment is automatically triggered (for example through residual activation or implicit learning; Chang et al., 2006; Pickering and Branigan, 1998) whenever a speaker is exposed to an expression with particular linguistic features.

But the greater alignment to 'computer' than 'human' interlocutors when the verb was repeated suggests the influence of an additional mediated component that is sensitive to beliefs about the interlocutor. In this case, such a component appears to be primarily associated with communicative enhancement and not social affect. Alignment based on social factors such as reciprocity or politeness should have occurred in the same way in the 'computer' condition as in the 'human' condition if the 'computer' were treated as a social agent comparable to the 'human', or to a lesser extent if it were treated as an inferior social agent. But Branigan et al. (2003) found *more* alignment in the 'computer' condition. This is most straightforwardly explained in terms of a strategy to enhance communication according to beliefs about the ability of the interlocutor.

Such a strategy does not appear to be wholly implicit, as stronger alignment with a 'computer' was only found when the verb was repeated. Repetition of the verb, together with the use of typed responses in which their utterance was visible on-screen, may have made participants more aware of the differences between the PO and DO constructions, allowing participants to choose to align or not. So it is possible that beliefs about an interlocutor affect alignment only when speakers are aware (to at least some extent) that a strategy of alignment is available. Branigan et al.'s (2003) results therefore suggest that alignment in HCI contexts may be influenced by beliefs about interlocutors' ability and a desire to enhance communicative success; if there is an effect of social affect, then it must be much smaller by comparison.

Converging evidence comes from a further study that investigated lexical alignment. Branigan et al. (2004) used the same experimental paradigm as Branigan et al. (2003), but this time participants named single objects rather than events. These objects had one highly preferred name (e.g., *bench*) and one highly dispreferred but acceptable name (e.g., *seat*), established by a pre-test. In one experiment, participants received either the highly preferred name or the highly dispreferred but

acceptable name from their interlocutor, and then had to name the same object for their interlocutor on a subsequent turn, as in the following interaction:

Interlocutor: <i>seat</i>	[Target object, dispreferred name]
Participant: <i>owl</i>	[Filler]
Interlocutor: <i>button</i>	[Filler]
Participant: <i>seat</i>	[Target object, dispreferred name]

In a further experiment, participants had to name the object on two turns (and as expected always used the preferred name) before receiving either the highly preferred or the highly dispreferred name from their interlocutor; they then had to name the same object again. By initially naming the object twice, participants established a strong precedent for both interlocutors to use that name within the interaction; for the interlocutor to use a different term could therefore be construed as impolite. Under these circumstances, one might expect participants to be less likely to subsequently align with their interlocutor, particularly on the use of a highly dispreferred term. In both experiments, Branigan et al. examined whether participants used the same name as their interlocutor or the alternative name.

Both experiments found lexical alignment to both 'computer' and 'human' interlocutors. However, participants were considerably more likely to align when they believed that they were interacting with a computer. For example, in the second experiment, participants broke their own precedent and aligned on a dispreferred term 67% of the time when interacting with a 'computer' but only 15% of the time when interacting with a 'human'. These results demonstrate that lexical alignment occurs in HCI. They also provide evidence that – like syntactic alignment – such alignment is influenced by beliefs about an interlocutor.

Note that these beliefs were independent of the interlocutor's actual behavior, as the 'computer' and 'human' interlocutor displayed identical behavior throughout. Moreover, this behavior consistently demonstrated to the participant that communication was always occurring successfully. Participants always received visual feedback on-screen that their interlocutor had correctly understood their description; for example, if they had typed *bench*, they would see a visual confirmation that the 'computer' or 'human' interlocutor had correctly chosen that object in response to their description. In particular, when participants had to name an object for the third time in the second experiment, they had already twice received evidence that the interlocutor understood their original choice of name. Yet they still overwhelmingly repeated the name that their interlocutor had just used, even if it was highly dispreferred, when they believed they were interacting with a computer.

From this finding we can conclude that alignment can be a very strong influence in HCI, overriding both strong linguistic preferences, and reciprocity and politeness norms. The finding that speakers aligned to a considerably greater degree when they believed they were interacting with a computer than with a person converges with Branigan et al.'s (2003) results in suggesting that social affect plays a relatively weak role in alignment in HCI and that considerations of communicative success play a larger role. Furthermore, the beliefs that mediated alignment in Branigan et al.'s (2004) study appeared to be based on a priori assumptions about the likely capability of a computer, and were not apparently updated during the interaction on the basis of direct evidence. This contrasts with evidence from dialogue between humans, which suggests that speakers are highly sensitive to the feedback that they receive from their interlocutors and update their model of the interlocutor dynamically. For example, Pearson et al. (2006b) used the same experimental paradigm as Branigan et al. (2004), but told participants that their interlocutor was another person. In one condition, they were told that their interlocutor was a native speaker of English; in the other, they were told that their interlocutor was a non-native speaker. Pearson et al. found that when participants received no feedback about communicative success, they showed more lexical alignment with a 'non-native' than with a 'native' interlocutor. This is consistent with them making a priori assumptions about the likely (limited) linguistic capability of a non-native speaker. But when participants received concurrent feedback that communication was successful, and specifically that their own initial choice of name for an object had been understood (exactly as in Branigan et al., 2004), they showed the same level of alignment with a 'non-native' speaker as with a 'native' speaker. This suggests that they used feedback dynamically to update their model of their human interlocutor's capability.

6.2.3. What can affect the magnitude of alignment in HCI?

But what factors affect the setting of such a priori expectations in HCI? We noted in Section 3.2 that in interactions between people, speakers draw a variety of inferences about their interlocutor based on assumptions about community membership, which may themselves be based on features such as appearance, accent and so on. Do analogous factors affect people's beliefs about computers in a way that would influence their propensity to align with them? Pearson et al. (2006a) investigated this question using the same method as Branigan et al. (2004). In their study participants were always led to believe that they were interacting with a computer (i.e., there was no human-interlocutor condition). Pearson et al. examined whether participants would generate different expectations about the computer on the basis of aspects of the interface design that were unrelated to its actual functioning. When participants arrived for the experiment, they were asked to wait while the system booted up. During this time, they saw a start-up screen. Participants saw one of two versions: a 'Basic' version that displayed the term *Basic version*, bore a 1987-dated copyright, and displayed a fictional computer magazine review stressing its limited features but cheap price and value for money; or an 'Advanced' version that displayed

the term *Advanced version: Professional edition*, bore a current-year copyright, and displayed a fictional computer magazine review stressing its expense and its impressive range of features and sophisticated technology. This start-up screen was the only difference between the 'Basic' and 'Advanced' conditions. In all other respects, participants in the 'Basic' and 'Advanced' conditions had identical experiences and the computer behaved identically during the actual interaction. Pearson et al. were concerned with whether such non-functional manipulations (i.e., manipulations that did not affect what the computer could do or how it functioned) would alter participants' beliefs about the system and its capability, and whether this would affect the way in which they interacted with it with respect to lexical alignment.

Despite experiencing identical behavior from the computer in both conditions, participants subsequently rated the 'Advanced' version as significantly more competent than the 'Basic' version. So an apparently trivial difference in the way in which a computer was presented significantly affected people's beliefs about it. But more importantly, these differences in beliefs had a significant effect on the way in which people interacted with it: Participants lexically aligned with both the 'Basic' and the 'Advanced' versions, but they aligned more strongly with the 'Basic' version. This finding is consistent with Branigan et al.'s (2003, 2004) results in suggesting that alignment in HCI is mediated by beliefs about the interlocutor's capabilities, and that the strongest component of alignment in HCI appears to relate to considerations of communicative enhancement. The results further show that such beliefs can be established by superficial non-functional properties of the computer that do not relate to its actual performance. Furthermore, as in Branigan et al. (2003, 2004), once established, these beliefs are not readily modified: Despite experiencing identical – error-free – behavior from the computer in both conditions, participants persisted in adapting their lexical choices more strongly to fit those of a computer that they had been led to believe had more restricted abilities.

7. How does alignment in HCI compare with alignment in human interactions?

It is clear that there is now strong evidence that linguistic alignment occurs in interactions between humans and computers, and that it takes place at many levels of structure. Indeed, we have seen that comparisons of alignment in HCI with alignment in interactions between humans consistently suggest that alignment is actually stronger in HCI than in interactions between humans.

The studies reviewed above suggest that the same mechanisms may underlie alignment in HCI as in interactions between humans, but that they may make different contributions. It is highly likely that there is an unmediated component that is activated merely by exposure to a particular linguistic structure and that is dependent on automatic priming mechanisms. But this component may be less important in HCI than it is hypothesized to be in human interactions (Pickering and Garrod, 2004). The evidence we have reviewed suggests strongly that interactions with computers are heavily influenced by considerations of communicative success and how to achieve it, and that these considerations are a potent determinant of alignment effects in HCI. It is striking how strong such effects are: For example, they can override extremely strong linguistic preferences, by causing people to use a name for an object that more than 85% of the time they would not otherwise use (Branigan et al., 2004).

The beliefs about capability that appear to largely underlie alignment in HCI appear to be driven to a great extent by the speaker's initial model of the computer interlocutor and what it is likely to know. As Pearson et al. (2006a) showed, this model may be based upon very superficial factors. In some ways this is comparable to speakers establishing beliefs about the community membership of human interlocutors based on their appearance. But the difference is that whereas in human interactions such beliefs are rapidly updated on the basis of the interlocutor's actual behavior, Branigan and colleagues' research suggests that people adopt a very conservative model of their interlocutor in HCI (Branigan et al., 2003, 2004; Pearson et al., 2006a). Even if they received consistent evidence of communicative success, and even evidence that a particular expression had been understood correctly, participants in their experiments nevertheless aligned with the most recent (and otherwise highly dispreferred) expression that their 'computer' interlocutor produced. Other evidence suggests some rapid adaptation, for example in the acoustic-prosodic properties of utterances (e.g., Oviatt et al., 2004). But these levels of alignment seem less likely to be subserved by a component concerned with enhancing communicative success.

One context of human interaction that is more closely analogous to HCI, and where we might therefore expect to see similar patterns of effects, is interactions between native and non-native speakers. In these cases, a native speaker may have similar considerations of how to ensure that communication is successful, given that her interlocutor may have different abilities to her own. We might therefore expect more alignment by a native speaker to a non-native speaker than to a native speaker, just as we find more alignment with a computer than with a human in HCI. Here the evidence seems mixed. Bortfeld and Brennan (1997) and Pearson et al. (2006b) found that native speakers aligned with non-native speakers (or in Pearson et al.'s case, with an interlocutor that they believed was a non-native speaker) on highly dispreferred (or erroneous) words, as speakers did with computer interlocutors in HCI (Branigan et al., 2004; Pearson et al., 2006a). But whereas Branigan et al. (2003, 2004) found that speakers showed more lexical alignment with computers than with humans irrespective of feedback, Bortfeld and Brennan (1997) found that native speakers did not align more overall with non-native speakers than with native speakers, and Pearson et al. (2006b) found that speakers showed more lexical alignment with a 'non-native' interlocutor than with a 'native' interlocutor only when they received no feedback; when they received feedback that communication was successful, they aligned to the same extent in both conditions.

Further research is needed to establish the precise commonalities and differences between HCI and native/non-native speaker interactions. A potentially critical dimension along which they may differ concerns the nature of the initial model

that speakers can establish in each case, and the way in which this is updated. When a native speaker encounters a non-native speaker, she may not know his level of linguistic competence, but she can make certain assumptions about his overall cognitive competence and his likely knowledge of the world. If interaction with him shows that these assumptions are in some way wrong, she can easily update them on the basis of her knowledge about the social distribution of linguistic and other knowledge (Fussell and Krauss, 1992). But she will have a good model of what the members of different communities are likely to know. With respect to speech communities, she can probably make a fairly good estimate that if he knows expression X, he is also likely to know expression Y, or that if he understands expression X, then he is likely to understand expression Y. This is because natural language acquisition generally follows a predictable pattern (and people are aware of that pattern). From a limited amount of evidence, then, speakers can form a relatively accurate model of their listener.

With computers, the position is completely different. Speakers are unlikely to have a good model of what computers are likely to know, presumably in part because natural language systems tend to be tailored to particular and quite specific uses. A particular system may have an extensive vocabulary related to an esoteric domain (e.g., financial instruments) but not know the most frequently occurring nouns in that language. In this sense, it is not possible to predict what kinds of expression a computer is or is not likely to know on the basis of limited experience in the way that it is possible to when interacting with a non-native speaker. That is, whereas any person can make a good guess to which speech community a person belongs, and what members of that community are likely to understand, it is not clear that people can make those guesses for a computer with any degree of accuracy.

At the beginning of this paper we noted that alignment may be critical to successful communication, and posited that HCI might be inherently doomed to communicative failure, if it turned out that alignment does not occur in HCI. The evidence that we have discussed suggests that it is not doomed in that respect: Alignment is often actually stronger when people are interacting with computers than with other people. But its strength depends upon speakers' beliefs about how to achieve successful communication with computers, which in turn depends on beliefs about their capability. We hypothesize that people may not feel confident in their ability to make accurate judgements about the capability of a computer, and might therefore tend to adopt a conservative initial model of the computer. Having adopted such a model, they might be reluctant to update it dynamically, or they might require substantial evidence of the computer's abilities before updating it. If our hypothesis is correct, then patterns of alignment in HCI may differ fundamentally from those found in interactions between humans, until increasing experience with technology or improvements in the technology mean that people are more confident in their assessments of computers' abilities. In other words, until we can adequately model a computer's knowledge in the way that we can model another person's knowledge, we may never interact with them in the same way that we interact with another person.

8. Conclusions

Research from many different perspectives over the last twenty years has converged in identifying the importance of behavioral alignment in human interaction. Such alignment occurs at non-verbal levels, but is particularly marked with respect to linguistic structure. Linguistic alignment in interactions between humans may play a critical role in achieving successful communication, and may arise from automatic priming processes that are encapsulated within the linguistic system, from social affective considerations, from considerations of communicative success, or from any combination of the three. Their relative contributions may vary according to the precise communicative context. There is now growing evidence that the same processes are operative in interactions between humans and computers, and indeed that they occur to a greater extent. In HCI, alignment appears to have a stronger strategic component that aims to maximise communicative success, reflecting speakers' beliefs about the limited capabilities of computers and the ensuing risk of communicative failure. Alignment by computers with human users can also induce positive affect in those users. Investigating linguistic alignment in HCI and comparing it with alignment in interactions between humans may cast light on the nature of human interaction in general, and may also provide critical behavioral evidence that can be leveraged to develop more effective human–computer interfaces.

Acknowledgements

This research was supported by a Scottish Enterprise grant (Edinburgh–Stanford Link), by ESRC grant RES-062-23-0376, and by a British Academy Postdoctoral Fellowship.

References

- Bailenson, Jeremy N., Yee, Nick, 2005. Digital chameleons—automatic assimilation of nonverbal gestures in immersive virtual environments. *Psychological Science* 16 (10), 814–819.
- Bavelas, Janet B., Black, Alex, Lemery, Charles R., Mullett, Jennifer, 1986. "I show how you feel": motor mimicry as a communicative act. *Journal of Personality and Social Psychology* 50 (2), 322–329.
- Bell, Allan, 1984. Language style as audience design. *Language in Society* 13 (2), 145–204.
- Bell, Linda, Gustafson, Joakim, Heldner, Mattias, 2003. Prosodic adaptation in human–computer interaction. *Proceedings of the International Congress of Phonetic Sciences (CD-ROM)* 2453–2456.
- Bortfeld, Heather, Brennan, Susan E., 1997. Use and acquisition of idiomatic expressions in referring by native and non-native speakers. *Discourse Processes* 23, 119–147.

- Bradac, James J., Mulac, Anthony, House, Ann, 1988. Lexical diversity and magnitude of convergent versus divergent style shifting perceptual and evaluative consequences. *Language and Communication* 8 (3–4), 213–228.
- Branigan, Holly P., Pickering, Martin J., Cleland, Alexandra A., 2000. Syntactic coordination in dialogue. *Cognition* 75 (2), B13–B25.
- Branigan, Holly P., Pickering, Martin J., Pearson, Jamie, McLean, Janet F., Nass, Clifford I., 2003. Syntactic alignment between computers and people: the role of belief about mental states. In: Alterman, David, Kirsh, Richard (Eds.), *Proceedings of the Twenty-fifth Annual Conference of the Cognitive Science Society*, Erlbaum, Mahwah, NJ, pp. 186–191.
- Branigan, Holly P., Pickering, Martin J., McLean, Janet F., 2005. Priming prepositional-phrase attachment during comprehension. *Journal of Experimental Psychology: Learning, Memory and Cognition* 31 (3), 468–481.
- Branigan, Holly P., Pickering, Martin J., Pearson, Jamie, McLean, Janet F., Nass, Clifford I., Hu, John, 2004. Beliefs about mental states in lexical and syntactic alignment: Evidence from human–computer dialogs. Paper Presented at the 17th Annual CUNY Human Sentence Processing Conference.
- Brennan, Susan E., 1991. Conversation with and through computers. *User Modeling and User-Adapted Interaction* 1, 67–86.
- Brennan, Susan E., 1996. Lexical entrainment in spontaneous dialog. *Proceedings of the International Symposium on Spoken Dialogue*, pp. 41–44.
- Brennan, Susan E., Clark, Herbert H., 1996. Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory and Cognition* 22 (6), 1482–1493.
- Chang, Franklin, Dell, Gary S., Bock, Kathryn, 2006. Becoming syntactic. *Psychological Review* 113 (2), 234–272.
- Chartrand, Tanya L., Bargh, John A., 1999. The chameleon effect: the perception-behavior link and social interaction. *Journal of Personality and Social Psychology* 76 (6), 893–910.
- Clark, Herbert H., 1996. *Using Language*. Cambridge University Press, Cambridge.
- Cleland, Alexandra A., Pickering, Martin J., 2003. The use of lexical and syntactic information in language production: evidence from the priming of noun phrase structure. *Journal of Memory and Language* 49 (2), 214–230.
- Costa, Albert, Pickering, Martin J., Sorace, Antonella, 2008. Alignment in second language dialogue. *Language and Cognitive Processes* 23 (4), 528–556.
- Danet, Brenda, 1980. ‘Baby’ or ‘Fetus’? Language and the construction of reality in a manslaughter trial. *Semiotica* 32 (3–4), 187–219.
- Fadiga, Luciano, Craighero, Laili, Buccino, Giovanni, Rizolatti, Giacomo, 2002. Speech listening specifically modulates the excitability of tongue muscles: a TMS study. *European Journal of Neuroscience* 15 (2), 399–402.
- Ferguson, Charles, 1975. Toward a characterization of English foreigner talk. *Anthropological Linguistics* 17 (1), 1–14.
- Fogg, B.J., Nass, Clifford I., 1997. Silicon sycophants: the effects of computers that flatter. *International Journal of Human–Computer Studies* 46 (5), 551–561.
- Fowler, Carol A., Brown, Julie M., Sabadini, Laura, Weihs, Jeffrey, 2003. Rapid access to speech gestures in perception: evidence from choice and simple response time tasks. *Journal of Memory and Language* 49 (3), 396–413.
- Fussell, Susan E., Krauss, Robert M., 1992. Coordination of knowledge in communication: effects of speakers’ assumptions about what others know. *Journal of Personality and Social Psychology* 62 (3), 378–391.
- Garrod, Simon, Anderson, Anthony, 1987. Saying what you mean in dialogue: a study in conceptual and semantic co-ordination. *Cognition* 27 (2), 181–218.
- Garrod, Simon, Clark, Aileen, 1993. The development of dialogue co-ordination skills in schoolchildren. *Language and Cognitive Processes* 8 (1), 101–126.
- Giles, Howard, Coupland, Nikolas, Coupland, Justine, 1991. Accommodation theory: communication, context, and consequence. In: Giles, Howard, Coupland, Justine, Coupland, Nikolas (Eds.), *Contexts of Accommodation: Developments in Applied Sociolinguistics*. Cambridge University Press, Cambridge, pp. 1–68.
- Giles, Howard, Powesland, Peter F., 1975. *Speech Style and Social Evaluation*. Academic Press, London.
- Goldinger, Stephen D., 1998. Echoes of echoes? An episodic theory of lexical access. *Psychological Review* 105 (2), 251–279.
- Gouldner, Alvin W., 1960. The norm of reciprocity: a preliminary statement. *American Sociological Review* 25, 161–178.
- Gries, Stefan Th., 2005. Syntactic priming: a corpus-based approach. *Journal of Psycholinguistic Research* 34 (4), 365–399.
- Hartsuiker, Robert J., Pickering, Martin J., Veltkamp, Elena, 2004. Is syntax separate or shared between languages? Cross-linguistic syntactic priming in Spanish/English bilinguals. *Psychological Science* 15 (6), 409–414.
- Hatfield, Elaine, Cacioppo, John T., Rapson, Richard L., 1994. *Emotional Contagion*. Cambridge University Press, New York.
- Haywood, Sarah L., Pickering, Martin J., Branigan, Holly P., 2005. Do speakers avoid ambiguities during dialogue? *Psychological Science* 16 (5), 362–366.
- Horton, W.S., Gerrig, R.J., 2005. The impact of memory demands on audience design during language production. *Cognition* 96 (2), 127–142.
- Horton, William S., Keysar, Boaz, 1996. When do speakers take into account common ground? *Cognition* 59 (1), 91–117.
- Isaacs, Eileen A., Clark, Herbert H., 1987. References in conversations between experts and novices. *Journal of Experimental Psychology: General* 116 (1), 26–37.
- Lakin, Jessica L., Chartrand, Tanya L., 2003. Using nonconscious behavioral mimicry to create affiliation and rapport. *Psychological Science* 14 (4), 334–339.
- Lakin, Jessica L., Jefferis, Valerie E., Cheng, C.M., 2003. The chameleon effect as social glue: evidence for the evolutionary significance of nonconscious mimicry. *Journal of Nonverbal Behavior* 27 (3), 145–162.
- Levelt, Willem J.M., Kelter, Stephanie, 1982. Surface form and memory in question answering. *Cognitive Psychology* 14 (1), 78–106.
- Linell, Per, 1998. *Approaching Dialogue: Talk, Interaction, and Contexts in a Dialogical Perspective*. Benjamins, Amsterdam.
- Maurer, Richard E., Tindall, Jeffrey H., 1983. Effect of postural congruence on client’s perception of counselor empathy. *Journal of Counseling Psychology* 30 (2), 158–163.
- Metzing, Charles, Brennan, Susan E., 2003. When conceptual pacts are broken: Partner-specific effects in the comprehension of referring expressions. *Journal of Memory and Language* 49 (2), 201–213.
- Meyer, David E., Schvaneveldt, Roger W., 1971. Facilitation in recognizing pairs of words: evidence of a dependence between retrieval operations. *Journal of Experimental Psychology* 90 (2), 227–234.
- Moon, Youngme, Nass, Clifford I., 1996. How “real” are computer personalities? Psychological responses to personality types in human–computer interaction. *Communication Research* 23 (6), 651–674.
- Nass, Clifford I., Fogg, B.J., Moon, Youngme, 1996. Can computers be teammates? *International Journal of Human–Computer Studies* 45 (6), 669–678.
- Nass, Clifford I., Lee, Kwan Min N., 2001. Does computer-synthesized speech manifest personality? Experimental tests of recognition, similarity-attraction, and consistency-attraction. *Journal of Experimental Psychology: Applied* 7 (3), 171–181.
- Nass, Clifford I., Moon, Youngme, 2000. Machines and mindlessness: social responses to computers. *Journal of Social Issues* 56 (1), 81–103.
- Nass, Clifford I., Moon, Youngme, Carney, Paul, 1999. Are people polite to computers? Responses to computer-based interviewing systems. *Journal of Applied Social Psychology* 29 (5), 1093–1110.
- Nass, Clifford I., Moon, Youngme, Fogg, B.J., Reeves, Byron, Dryer, Christopher, 1995. Can computer personalities be human personalities. *International Journal of Human–Computer Studies* 43 (2), 223–239.
- Oviatt, Sharon L., Bernard, Jon, Levow, Gina-Anne, 1998. Linguistic adaptations during spoken and multimodal error resolution. *Language and Speech* 41 (3–4), 419–442.
- Oviatt, Sharon L., Darves, Courtney, Coulston, Rachel, 2004. Toward adaptive conversational interfaces: modeling speech convergence with animated personas. *ACM Transactions on Computer–Human Interaction* 11 (3), 300–328.
- Pardo, Jennifer S., 2006. On phonetic convergence during conversational interaction. *Journal of the Acoustical Society of America* 119 (4), 2382–2393.
- Pearson, Jamie, Hu, John, Branigan, Holly P., Pickering, Martin J., Nass, Clifford I., 2006a. Adaptive language behavior in HCI: how expectations and beliefs about a system affect users’ word choice. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM Digital Library, pp. 1177–1180.
- Pearson, Jamie, Pickering, Martin J., Branigan, Holly P., Nass, Clifford I., Hu, John, 2006b. Influence of prior beliefs and (lack of) evidence of understanding on lexical alignment. Paper Presented at the Architectures and Mechanisms of Language Processing Conference.
- Pickering, Martin J., Branigan, Holly P., 1998. The representation of verbs: evidence from syntactic priming in language production. *Journal of Memory and Language* 39 (4), 633–651.

- Pickering, Martin J., Ferreira, Victor S., 2008. Structural priming: a critical review. *Psychological Bulletin* 134 (3), 427–459.
- Pickering, Martin J., Branigan, Holly P., McLean, Janet F., 2002. Constituent structure is formulated in one stage. *Journal of Memory and Language* 46 (3), 586–605.
- Pickering, Martin J., Garrod, Simon, 2004. Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences* 27 (2), 169–225.
- Reeves, Byron, Nass, Clifford I., 1996. *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. Cambridge University Press, New York.
- Schober, Michael, 1993. Spatial perspective-taking in conversation. *Cognition* 47 (1), 1–24.
- Suzuki, Noriko, Katagiri, Yasuhiro, 2007. Prosodic alignment in human–computer interaction. *Connection Science* 19 (2), 131–141.
- Tannen, Deborah L., 1987. Repetition in conversation: toward a poetics of talk. *Language* 63 (3), 574–605.
- van Baaren, Rick B., Holland, Rob W., Steenaert, Bregje, van Knippenberg, Ad, 2003. Mimicry for money: behavioral consequences of imitation. *Journal of Experimental Social Psychology* 39 (4), 393–398.
- van Baaren, Rick B., Holland, Rob W., Kawakami, Kerry, van Knippenberg, Ad, 2004. Mimicry and prosocial behavior. *Psychological Science* 15 (1), 71–74.
- Ward, Nigel, Nakagawa, Satoshi, 2002. Automatic user-adaptive speaking rate selection for information delivery. *Proceedings of the International Conference on Spoken Language Processing*, pp. 549–552. www.isca-speech.org/archive/icslp_2002/.
- Watkins, Kate, Paus, Tomáš, 2004. Modulation of motor excitability during speech perception: the role of Broca's area. *Journal of Cognitive Neuroscience* 16 (6), 978–987.
- Watson, Matthew E., Pickering, Martin J., Branigan, Holly P., 2004. Alignment of reference frames in dialogue. *Proceedings of the 26th Annual Conference of the Cognitive Science Society*, Chicago, Erlbaum, Mahwah, NJ, pp. 2353–2358.

Holly Branigan is Reader in the Department of Psychology at the University of Edinburgh. She holds a BA in Language and Linguistic Science from the University of York, and an MSc and PhD in Cognitive Science from the University of Edinburgh. Her research focuses on language production in adults and children, with a particular interest in syntactic processing and representation; other research areas include language use in dialogue, human–computer interaction, and dyslexia.

Martin Pickering is Professor of the Psychology of Language and Communication in the Department of Psychology at the University of Edinburgh. He holds a PhD in Cognitive Science from the University of Edinburgh and is a Fellow of the Royal Society of Edinburgh. He has published extensively in language comprehension, language production, dialogue, bilingualism, and the relationship between psychology and linguistics.

Jamie Pearson is a senior information analyst and statistician at the Information Services Division, NHS National Services Scotland. He earned his BA in Psychology at the University of Liverpool and his PhD in psycholinguistics at Durham University.

Janet Mclean is a research fellow of the Department of Psychology at the University of Edinburgh. She earned a BSc in Maths and Psychology at the Newcastle University, and a PhD in Psychology from Lancaster University. Her research focuses on dialogue and language production in both children and adults.