Meaning Negotiation

Massimo Warglien, Ca Foscari University, Venezia
Peter Gärdenfors, Lund University

Abstract

While "meaning negotiation" has become an ubiquitous term, its use is often confusing. A negotiation problem implies not only a convenience to agree, but also diverging interest on what to agree upon. It implies agreement but also the possibility of (voluntary) disagreement. In this chapter, we look at meaning negotiation as the process through which agents starting from different preferred conceptual representations of an object, an event or a more complex entity, converge to an agreement through some communication medium. We shortly sketch the outline of a geometric view of meaning negotiation, based on conceptual spaces. We show that such view can inherit important structural elements from game theoretic models of bargaining – in particular, in the case when the protagonists have overlapping negotiation regions, we emphasize a parallel to the Nash solution in cooperative game theory. When acceptable solution regions of the protagonists are disjoint, we present several types of processes: changes in the salience of dimensions, dimensional projections and metaphorical space transformations. None of the latter processes are motivated by normative or rationality considerations, but presented as argumentation tools that we believe are used in actual situations of conceptual disagreement.

1. Introduction

"Meaning negotiation" has become an ubiquitous term, used in contexts as diverse as semantics and epistemology (Larson and Ludlow 1993), conversation theory (Brennan an Clark1996), ethnography (Wenger 1998), but also literary criticism, artificial intelligence, psychotherapy. The concept suggests that meaning is often not uniquely determined by the lexicon and ordinary utterances, and thus there is room left for a process of further determination through some type of interaction among communicating agents.

However, the expression "meaning negotiation" may be sometimes a source of confusion. In many current usages of the term, a negotiation is often confused with an agreement. However, a negotiation problem implies not only a convenience to agree, but also diverging interest on what to agree upon. It implies agreement but also the possibility of (voluntary) disagreement. Thus, a problem of negotiation differs from a problem of pure coordination, since while negotiators both have an interest to agree (as in coordination), they nevertheless have conflictual interests in dividing the value generated by their cooperation. Schelling (1960) has described this type of interaction as one of "mixed motives", since common interest and conflict coexist in the same situation.

Robert Stalnaker has well captured the issue of "mixed motives" in his description of a conversation game:

"*One may think of a non-defective conversation as a game where the common context set is the playing field and the moves are either attempts to reduce the size of the set in a certain ways or rejections of such moves by others. The participants have a common interest in reducing the size of the set, but their interests may diverge when it comes to the question of how it should be reduced.*" *(Stalnaker 1999)*

In this chapter, we will look at meaning negotiation as the process through which agents starting from different preferred conceptual representations of an object, an event or a more complex entity, converge to an agreement through some communication medium. The process is typically a sequence of offers and counter-offers that are accepted or rejected (as in

Clark's "contributions" (Clark and Schaefer 1989). The "solution" to the negotiation problem is the agreement reached (or the final disagreement). While this approach maintains a broad scope, it is important to stress that it assumes that agents "move" in a defined conceptual space and have potentially conflicting interests in the agreement to be reached.

There are many examples of rather ordinary communication contexts in which issues of meaning negotiation arise very naturally.

A simple and powerful example has been defined by Furnas et al. (1987) as the "vocabulary problem". Studying human-system interactions, Furnas et al. found that even in simple naming tasks individuals rarely agree *ex ante* on which word to use for referring to common objects or situations. They found that in such cases there are in general no perfect synonyms, and there is a low probability of *ex ante* lexical agreement between two different individuals – this difficulty of *ex ante* lexical agreement is also at the core of a popular computer interactive web game, the ESP game by von Ahn (2006)[1].

As individuals do not agree *ex ante* on the lexical choice, and differences in their preferred one may actually mark subtle differences in the way they conceptualize the situation at hand, how do they converge on a sufficiently agreed lexicon during a conversation or other types of communicative interaction? Brennan and Clark (1996), in their analysis of the "vocabulary problem", have submitted that this happens through "conceptual pacts" – temporary agreements about how the referent is conceptualized. Once such a "pact" is reached, individuals can repeatedly and confidently refer to an object with the same term – which translates into the familiar phenomenon of lexical entrainment (Garrod and Anderson, 1987; Brennan, 1996; Pickering and Garrod, 2004), i.e. the tendency of people to adopt the terms introduced by their interlocutor within a conversation.

Brennan and Clark notice some features of conceptual pacts which are worth reporting here. First of all, "conceptual pacts are established by speakers and addressees jointly" (Brennan and Clark 1996, 149). They are the result of an interactive process that may involve different rounds, lexical proposals and counterproposals, and may imply also disagreement. Furthermore, lexical pacts are specific to a given speaker-addressee pair. In other words they tend to reflect the specific relation between the two and the process through which an

---

[1] In the ESP game (so called because it encourages players to "think like each other"), two players randomly matched through the web have to find a common (agreed) label for an image. The game has become a prototype for the concept of "Games with a purpose", since human participants' playfulness is used to solve problems that are difficult to solve in automated ways - in this case image labeling (von Ahn 2006).

agreement has been reached – the same speaker may reach different pacts with different addressees. The emergence of conceptual pacts on the early stages of a conversation has been shown to be a good predictor of the overall cooperative success of communication (Reitter et al., 2007; Nenkova et al., 2008).

Another interesting example of continued negotiation of concepts, where payoffs are not just semantic, comes from Andersson (1994). He investigates how different meanings of "nature" are used and argued for by different social, political and cultural groups. For example, he documents the tensions in the meaning of "natural forest" between forest owners, environmentalists and government officials and their power struggles to establish their preferred meaning. The outcome of the negotiations will have economic, environmental and legal consequences.

The use of vague predicates in communication provides another neat example of the ubiquitousness of meaning negotiation. When a vague predicate is asserted in communication, this often corresponds to a move that proposes to restrict the range of its possible values. As Barker (2002, 2013) suggests, by stating that "Harrison Ford is a tall actor" a speaker suggests that any actor taller than Harrison Ford is tall as well – if the addressee accepts this statement, all actors taller than Harrison Ford will be automatically annexed as tall to the common ground of the conversation. A parallel statement that "Tom Cruise is not a tall actor" would introduce a new restriction on the interval of tall actors, narrowing the range of admissible standards of tallness for actors. Of course, some of these statements may be rejected by an addressee – for example by rejecting the assertion that "Johnny Depp is tall" the addressee would signal her refusal to concede that the standard of tallness falls below the 1.80 m. limit.

The fact that vague predicates are intrinsically underdetermined invites their renegotiation in the context of each specific conversation. Through negotiation, agents can reach an agreement that sufficiently restricts the vague area to satisfy the coordination needs of communication – or decide that they cannot agree. As such, meaning negotiation contributes to the flexibility of vague predicates, and makes them adaptable to different contexts.

Another interesting, and more subtle, case of meaning negotiation is related to indirect speech (Pinker, Nowak and Lee, 2008). Indirect speech often reveals the presence of conflicting preferences in communication, and the need of communicating agents to negotiate through language the understanding of their mutual relations. Why should people often blur their communicative intents by allusive expressions or euphemisms? Pinker and coauthors suggests that, among other motivations, indirect speech reveals uncertainty about the

intentions of the listener, and are often first moves in a series of language manoeuvers allowing to explore possible agreements without incurring the psychological (or sometimes material) cost of rejection. For example, a driver trying to bribe the policeman fining him might try a phrase such as "so maybe the best thing would be to take care of it here" (Pinker et al., 2008, 834), thus checking the honesty of the policeman without making an explicit offer that might lead to an accusation of bribery. The strategic nature of indirect speech is even more apparent when communication does not play only the role of transmitting information but also supports the negotiation of reciprocal relations between two persons (as in the case of an allusive sexual offer).

2. **Negotiating in conceptual spaces**

Conceptual spaces (Gärdenfors, 2000) provide a very natural framework for modelling meaning negotiation. In Warglien and Gärdenfors (2013) and Gärdenfors (2014) we develop an account of meaning as emerging from the interaction of different individual conceptual representations – as a "meeting of minds". By taking a radical departure from traditional semantics, we state that the meaning of an expression does not reside in the world or (solely) in the mental schemes of individual users, but rather emerges from the mappings between individual mental spaces that are established through communication. The fundamental role of a communicative act, in this view, is to try to bring about cognitive changes (van Benthem, 2008) by affecting others' states of mind.

The "meeting of minds" framework is couched in a geometric view where concepts are represented as convex regions of conceptual spaces, and the emergence of meaning is modelled as resulting from the mutual convergence of the positioning of each agent in the "product space" of their mental representations.

A simple example of such a process, that can be used as a more general metaphor of the emergence of meaning from interaction, is the achievement of joint attention in children's pointing (Bates, 1976; Brinck, 2004; Gärdenfors and Warglien, 2013). A meeting of minds occurs in pointing when child and mother perceive that, as a result of an original directional gesture, they are aligning their focus of attention on the same point in the surrounding physical space. When this convergence happens and is mutually recognized (e.g. through mutual gazes), the child's picture of what he is pointing out to the mother agrees with his understanding of what she is attending to (the same for the mother), and a sort of

communicative equilibrium point is established, which can be formally modelled as a fixed point in the mappings between the mental spaces of the interacting agents (Warglien and Gärdenfors, 2013).

The emergence of meaning in linguistic communication can be seen as a sort of generalized pointing process, in which language is used to drive the other's mind in a desired direction in her own mental space. A formal analysis of such communication processes shows that convexity of concepts plays an important role in ensuring that a "meetings of minds" exist. Other features of linguistic communication further support the existence of such points and the possibility to reach them. For example, the fact that the lexicon can express the categorization of an underlying conceptual space allows the use of discrete language tokens to approximate fixed points, while pragmatic maxims of conversation à la Grice (1989) facilitate convergence to such points (Warglien and Gärdenfors 2013).

However, in Warglien and Gärdenfors's original formulation of interactive meaning, little attention has been paid to the role played by differences in individual preferences for a given conceptualization. For example, different lexical preferences applied to the same object may translate into different conceptualizations of that object. When two individuals referring to the same brick wall use "wall" and "barrier", they may categorize differently the same visual scene and express a different communicative intention, reflecting different preferences for the scene representation (they "point" to different conceptual entities). In this case, finding a mutual agreement may imply some lexical give and take through which a meaning negotiation happens.

In order to understand meaning negotiation, one needs to develop a notion of an individual commitment to some preferred representation – for example a given categorization of an object, a certain combination of quality features characterizing a product, or a representative example of a tall actor. In general, the nature of conceptual preferences can be purely cognitive – for example the result of the individual learning history. In other cases, it may reflect the value of a specific conceptualization in the light of broader utility considerations – for example, the interpretation of the prototypical quality of the object of exchange in a commercial contract (e.g. what "a workmanlike job" means in a construction contract) is subject to obvious conflict of interest between the two transacting parts. Furthermore, one needs to develop a notion of what makes it acceptable to diverge from such preferred representations, as well as of what will make divergence unacceptable – to the point that we might prefer an open disagreement.

A simple way to capture the essence of a meaning negotiation problem in conceptual spaces is to assume that individuals have preferred points in such space, and that there is a subjective cost in departing from a preferred point. For example, I may have my own preferred threshold for separating tall from non-tall actors in the dimension of height. I may be willing to accept departures from such threshold for the sake of conversation , but the larger the divergence from my own standard, the larger my discomfort. Thus, the cost of divergence from my favorite point will be a function of distance from such point (distance in conceptual spaces will express some measure of dissimilarity). And beyond a certain point, the discomfort caused by such divergence will offset the advantages of keeping our conversation running – and disagreement will break out.

The idea can be expressed graphically in a simple way. Consider a two-dimensional conceptual space (e.g. the space of beer "strength", defined by a combination of bitterness and alcohol degree). Let us assume that two individuals mutually engaged in communication share the same two-dimensional conceptual space (later we will relax this assumption), but have two different points defining a prototypical "strong beer", respectively A and B. While both have an interest in developing their conversation, assume that the expected benefits of a conversational agreement are reduced by any deviation from their preferred point. If the cost of such deviations is an increasing function of distance from the ideal point, after a given distance there may be no further interest in agreeing with the other agent, and conflict will be preferred to concession. Thus, for each agent, the area of acceptable definitions of a strong beer would be a circle (Fig. 1) around respectively point A and B. The circumference of the circles delimit what each player can afford as the maximum acceptable distance from the prototype. The intersection of such circles will define a set of possible agreements – what both can accept as a definition of strong beer (Fig. 1a). Not all the points of the possible agreement set have the same status, though. Points that are outside the segment connecting A and B are in a strong sense (Pareto) inefficient: agents can improve their position without damaging the other by coming closer to such segment. Thus, the bold segment (a,b) will define the efficient agreement set, where agreements should be expected to fall (Fig. 1b).
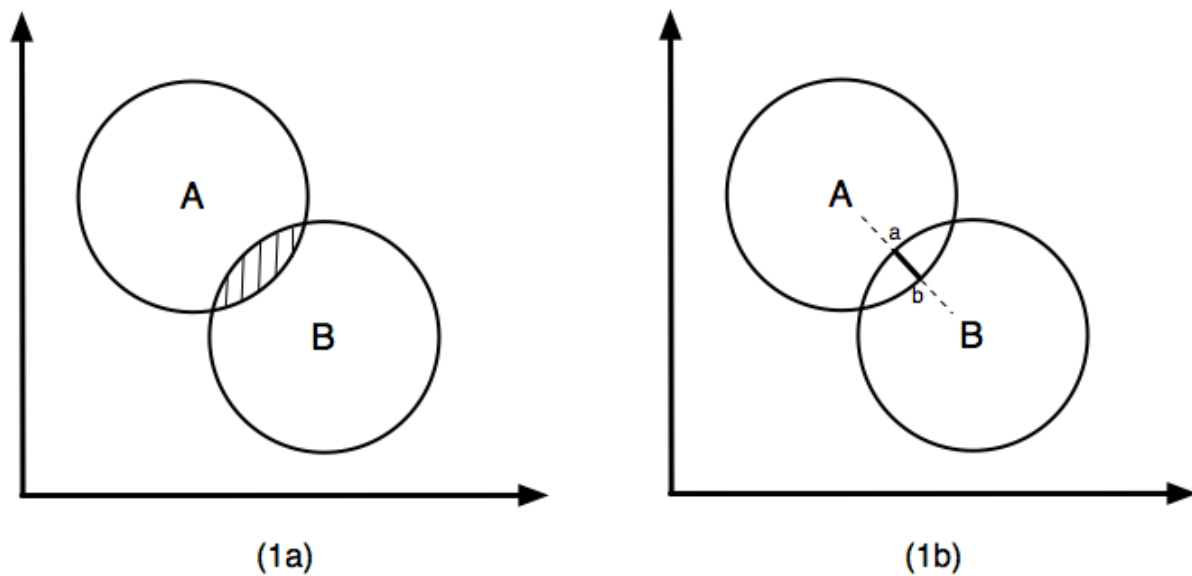
Figure 1 - Acceptable and effcient agreements

While this powerfully restricts the set of agreements one is expected to observe, it leaves still indetermined the problem of which agreement should prevail. Game theorists have provided a large repertoire of solution concepts for situations such as the one described above. One solution that most naturally fits communication contexts however pre-dates game theory – Zeuthen's (1930) approach to bilateral bargaining. Zeuthen's idea can be simply summarized by a key question: at each stage of a bargaining process, who should make a concession? The outcome of a bargaining problem will depend from the progressive contraction of the bargaining space as a result of subsequent concessions.

Imagine a situation like the one depicted in fig. 2. The possible agreements set is not empty, thus a solution to the meaning negotiation problem may be expected. Mary and Joe have both made initial proposals M1 and J1 that fall within the agreement set, but leave room for further negotiation. For example, Mary could accept Joe's offer (after all it's in the acceptable set), but of course her own proposal is more convenient to her. Its not clear, though, that Joe could accept Mary's proposal, so Mary could be tempted to make a concession by proposing, for example, M2. Similarly, Joe has to decide whether to stay on his own proposal – at the risk of triggering a conflict with Mary – or make in turn a concession, say J2.
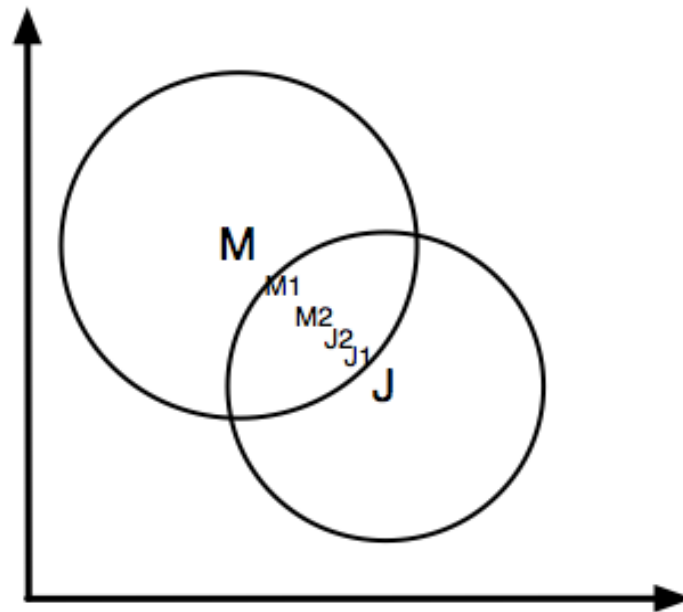
Figure 2 - Moves in the bargaining process

Zeuthen's key assumption is that at each turn Mary (Joe) will evaluate the situation, and assess what is the maximum acceptable risk that Joe (Mary) will prefer the conflict rather than accept M1 (J1). The intuition is that the one who has the lowest acceptable risk will concede. In other words, those that can afford a larger risk have an advantage. An important implication is that if agents are symmetric (same cost of conflict, same risk aversion) the solution will be an even split along the efficient agreements line. Otherwise, players enjoying a comparatively lower conflict cost and having higher propensity to risk will be more willing to engage in conflict and thus have higher bargaining power.

The idea that players more willing to face conflict have an advantage seems empirically reasonable also in the communication domain – and indeed experimental data on how individuals negotiate a language can support such claim. For example, Selten and Warglien's (2007) experiment on the emergence of a common (artificial) lexicon in a two-person language game of referring shows that players signaling their stubborness in the early stages of the game have a decisive advantage in imposing their own preferred code for referring to a set of objects. Brennan's (1996) study on lexical entrainment in man-machine interaction shows that individuals having a conversation with a computer dialogue interface tend to concede to the lexicon of the (credibly stubborn) machine.

Under the further assumption that individuals act according to expected utility, Zeuthen's mechanism leads to a well-known solution concept in cooperative game theory, the Nash bargaining solution (Harsanyi, 1956).[2] While we will not go into the details of such solution concept in this chapter, it is worth noting that the same topological properties (compactness and convexity) of conceptual spaces that support "meetings of minds" (Warglien and Gärdenfors, 2013) support the existence and uniqueness of the Nash solution.

This view of meaning negotiation crucially depends on the fact that some intial representation is established for each agent, and that agents can locate their meanings in such space. It may be questioned that it is always possible to open up new dimensions in a negotiation. However, since concepts have an "open texture" (Porosität) (Waismann 1968), there is always some new aspect of a concept that has not been captured by the negotiation. Waismann argues that concepts (outside mathematics) can never be given a complete definition in terms of necessary and sufficient conditions: "[T]here will always remain a possibiility ... that we have not taken into account something or other that may be relevant to their usage. ... I shall never reach a point where my description will be completed" (1968, p. 121-122). Broader situations in which meaning has to be jointly elaborated through a process of search of relevant dimensions (see Egré 2013) may lay outside the scope of our use of "meaning negotiation"

3. **Variations**

Different assumptions may lead to slightly different ways of determining the equilibrium solution (Thomson 1994) – and the fact that in many cases language is discrete may lead anyway to just approximations of such solutions (Warglien and Gärdenfors 2013). Equilibira solutions are motivated by normative considerations. However, there are also other argumentation tools that may be involved in reaching an agreement or a partial agreement in a negotiation. In this section we present some tools of this kind.

Interesting implications can be derived by assuming some form of bounded rationality, that may limit the ability of agents to have a full rational control of all the space of representations. For example, many conversational phenomena appear to be driven by

---

[2] The Nash solution predicts that players will jointy maximize the product of their utilities. The Nash solution shoud not be confounded with the concept of Nash equilibrium. The Nash equilibrium and the Nash solution to the bargaining problem belong to two different families of game theoretic solutions, the former being a non-cooperative games solution concept, the latter a cooperative games one: see Osborne and Rubinstein (1994) for an accessible introduction to both. For example, the Nash solution assumes Pareto-efficiency as an axiom, while Nash equilibria can be non Pareto-efficient.

automatic processes rather than deliberation (Pickering and Garrod 2004). Also, it has been shown that the choice of the reference word in lexical comparison can alter the salience of properties of given objects affecting similarity judgments (Ortony et al. 1985). In these cases, the effects of changes (or manipulations) of representations can be considered, that can substantially affect the outcome of meaning negotiation.

3.1 *Salience manipulation.*

Consider a case such as is depicted in figure 3. The distance (conceptual dissimilarity) between the preferred points C and D is such that it overrides the benefits of cooperation, and lack of agreement should be expected. However, the distance between such points depends on how the two dimensions are weighted – something that will depend crucially from the salience attributed on each dimension (Gärdenfors, 2000). Appropriate manipulation of the salience of different dimensions can modify the perception of dissimilarity between two agents (Ortony et al., 1985) and modify the distance between the two preferred points, thus facilitating the emergence of a possible agreement area (see figure 3). It is well known that salience effects can be manipulated in conversation and affect perceptions in automatic, hard-to-control ways (Taylor et al. 1979). For example, a speaker may exploit the priming effects of mentioning some words early in the conversation to make the dimensions associated to such words more relevant through entrainment, a mostly automatic process (Pickering and Garrod, 2004).
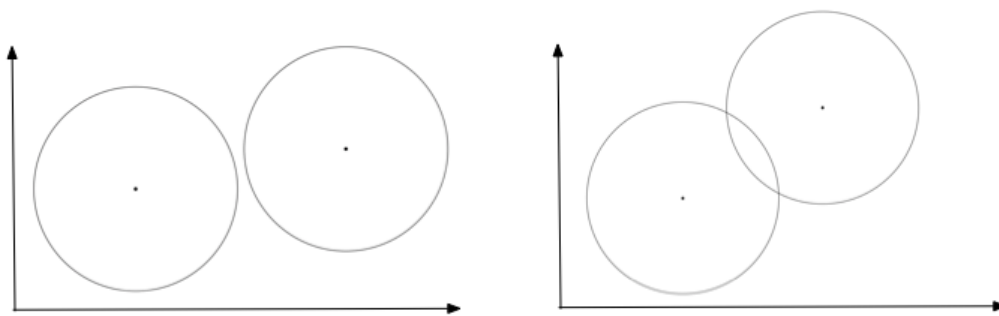

Figure 3 - Agreements are made possible as the weight of the abscissa is reduced

3.2 *Partial agreements*

12

The use of semantically underdeterminate words is frequent in language (Ludlow, to appear). While in some cases it can be related to simple reasons of economy, e.g. when a more determinate description is not needed, some level of indeterminacy may be related to the search for a partial agreement when a full one is not reachable. Typically, a partial agreement will consider only some dimensions of the problem, ignoring or deferring other ones. For example, as concepts can have multiple quality dimensions but adjectives typically represent only single or integral (non-separable) attributes (Gärdenfors 2000, 2014), a conversation focusing on a specific adjective will implicitly foster partial lexical pacts.
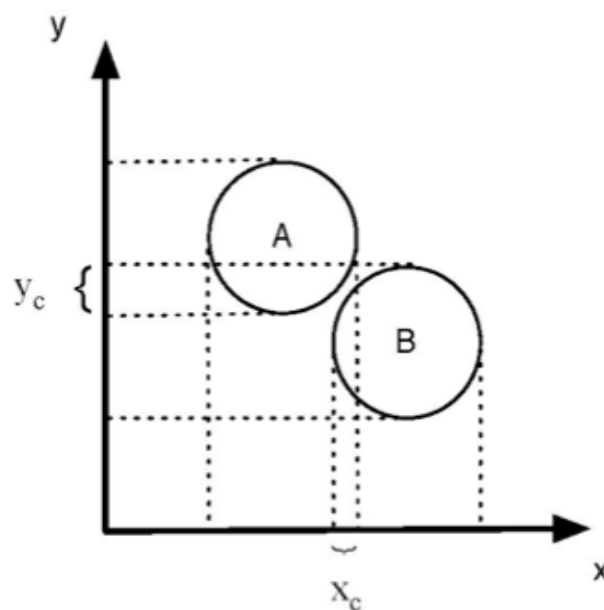

Figure 4 - Partial agreement areas on projections

Once more, conceptual spaces suggest a natural analysis of such a phenomenon. Consider again two individuals with apparently incompatible conceptualizations, such that the circles representing their area of acceptable meanings do not overlap (fig. 4). Despite the global incompatibility of their concepts, they still can agree on the projections of the circles on the single conceptual dimensions x or y (the x_c and y_c segments in fig. 4). Thus, there is room for at least a partial agreement on each dimension. One classical example is the separation of disagreement on facts from disagreement on values, the former being often easier to solve (Perelman and Olbrechts-Tyteca 1958/1969). On a micro-level, conversational templates such as the proverbial "It doesn't matter whether a cat is white or black, as long as it catches mice" provide an obvious template for partial agreements.

Of course, which dimension prevails might create some advantage for one of the two speakers, which implies that there is ample room for different forms of dialogue manipulation (Van Eemeren and Grootendorst, 2004; Van Eemeren, 2010) . For example, as the projection over a subset of conceptual dimensions can be considered as a degenerate case of salience manipulation (in which one or more dimensions have zero salience), all the conversational moves trying to focus on single properties (e.g. exploiting entrainment) can have the effect of facilitating a partial agreement on dimensions favorable to the speaker.

Partial agreements are also likely when representations of speakers have different dimensions, sharing only some of them. In that case, it appears as unavoidable that meaning negotiation will be performed on the shared dimensions, leaving others often implicit and thus leading to intrinsically underdetermined agreements.

While partial agreements emphasize dimensional reduction as a strategy to make an agreement possible, the symmetric manipulation, i.e. introducing new dimensions, is of course possible. Adding new dimensions might be motivated by the necessity to search for solutions in a broader space, but also respond to other strategic considerations, such as the need for an agent to move out of a negotiation space where he has a comparative disadvantage. We don't further elaborate this case here, although it is clearly relevant for meaning negotiation dynamics and it is a natural development of our approach.

3.3 *Metaphoric projection*

Students of negotiation often stress the key role that metaphors play in the linguistic interaction that lead to negotiated agreements. But how do metaphors affect meaning negotiation itself? We suggest that metaphors play a key role in meaning negotiation by performing at a same time a selection process over dimensions and a modification of the similarity structure of the discourse domain.

Metaphors are commonly understood as mappings that transfer structure from a source domain to a target one. Such mappings act selectively on both the source and target domains – they select specific structural aspects of the source and mold the target according to such structural aspects – thus only those dimensions of the target which are compatible with the target are selected: "the lion Ulysses" emphasizes Ulysses' courage but hides his condition of a castaway in Ogiya. Thus metaphors act by orienting communication and selecting

dimensions that may be more or less favorable to the speaker. By suggesting that a storm hit the financial markets, a bank manager can move the conversation away from dimensions pertaining to his own responsibilities and instead focus on dimensions over which he has no control, strengthening his position vis à vis his audience (Rocci, 2009).

At the same time, metaphors shape the distance between different points in conceptual spaces by providing context for their interpretation – e.g. by providing contrast classes within which distance between elements is modified. This can be illustrated by a more complex example, provided by how the Falling Dominoes metaphor, dominating foreign policy in the fifties and sixties, created a representation that brought close to each several countries otherwise differing in terms of political and military issues; for example, it downplayed those aspects of the North Vietnamese position related to nationalism to empasize ideological dimensions shared with other countries of the Communist block (McNamara et al. 2007). This led for example to significantly downplay the strong distance between North Vietnam and China. The Falling Dominoes metaphoric blindfold forced alternative positions in foreign policy into a funnel that significantly narrowed disagreements over possible policies.

3.4 *Non cooperative aspects of meaning negotiation*

Until now we have assumed that in meaning negotiation agents will agree on a point which is in the efficient set of possible agreements – the standard assumption of cooperative games. However, it may be useful to remove this assumption in order to analyze the emergence of conversational failures. LiCalzi and Maagli (2013) have analyzed the problem of negotiating the categorization of conceptual spaces using the tools of non-cooperative game theory. The example of the negotiations of the meaning of "natural forest" from Andersson (1994) is a real life illustration of such a situation.

There are two agents and each one of them has a conceptual space that for analytic convenience is represented as a circle. Each agent categorizes the conceptual space in two convex partitions, Left and Right (the dividing line needs not be a diameter). Agents have an incentive to agree on the same categorization of the circle, but if the agreed partition is different from their preferred one, they incur a disutility proportional to the area they are "giving up" to the agreement – so each rational player strives to minimize her losses.

Figure 5 - Two partitions

This simple structure allows singling out two types of initial disagreement. One is the "focused" disagreement, where the lines dividing the conceptual spaces for each agent do not cross each other – thus the disagreement area is a convex region between the two lines (fig. 6a) The other one is the "widespread disagreement", where the lines partitioning the conceptual spaces for each agent cross each other. This implies that the area of initial disagreement is not convex (fig. 6b). As we shall see, these initial conditions of disagreement have important implications for the solution reached.
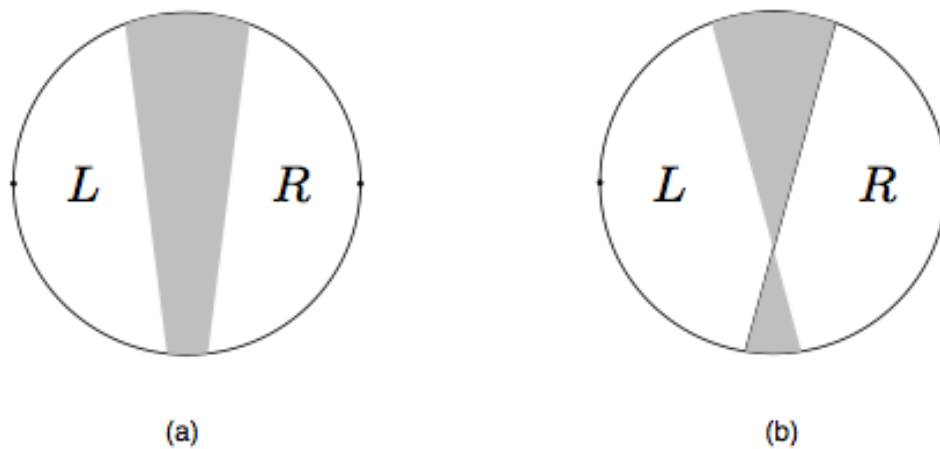

Figure 6 - Focused and widespread disagreement

The case of focused disagreement is simpler and illustrates how the game works. It basically happens to be a zero sum game in which what is lost by one player is gained by the other one. A simple sequential process gives the first move in the game to one player, mimicking the possibility of the first speaker to create an anchor to the establishment of a "common ground". He or she can thus choose where to locate a first point on the circle's circumference. The second player can only pick a second point that will determine how the circle is partitioned. If both players are trying to minimize their disutility, they will both concede

nothing of their original partition, and will stick to one extreme or their dividing line. The result is shown in fig. 7a: the player who controls the longest arc (e.g. the one who can control the communciation "agenda") loses less than the other (the shaded area represnts the initial disagrameent area).
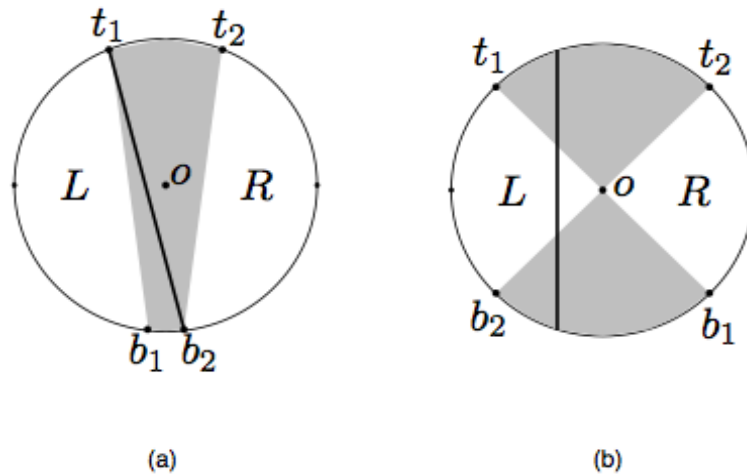

Figure 7 - Nash equilibria of the game for the focused and widespread disagreement

The case of widespread disagreement is more complex /fig 7b). To see it assume (with no loss of generality) that the two dividing lines (at the beginning of the game) go through the center of the circle. It would seem reasonable that at the end fo the game the dividing line passes through the center, providing an efficient solution. Unfortunately, the Nash equilibrum of the game, given the "stubborness "of players in minimizing individually their disutility, leads to a communication failure: both players lose some of the potential area of consensus, generating a solution that is inefficient — that could be improved if both acted in a more collaborative way. Thus, a simple conversation game can show the emergence to a sort of "conversational dilemma", the failure of communication to preserve the pre-existing consensus.

4. **Discussion**

Despite its pervasivness, meaning negotiation is still a rather under-analyzed penomenon. We can only speculate on why it is so. Three reasons stand as rather natural. First of all, meaning negotiaion presupposes a view of language in which semantic underdetermination plays an important role – a view certainly in contrast with the central tenets of classical semantics. Furthermore, it presupposes a view of meaning as (at least to some extent) a social,

interactive phenomenon, once more violating the strong view of meaning as fundamentally independent from the communicative interaction of speakers.

Finally, while the pragmatic tradition concedes a significant role to communicative interaction, it still relies heavily on the assumption that language in use is a collaborative enterprise (Clark 1996), leaving in the shadow aspects related to conflict between communicating agents.

We claim that a geometric approach to meaning (Gärdenfors; 2000, Warglien and Gärdenfors; 2013; Gärdenfors, 2014) is well equipped to deal with such "anomalous" features of meaning negotiation. It allows to explicitly represent underdeterminateness in terms of regions of a meaning space. It allows to naturally represent the interactive nature of meaning via mappings between different individual meaning spaces. And it can represent conflicting preferences for meanings as different locations in such spaces.

In this chapter we have shortly sketched the outline of a geometric view of meaning negotiation. We have shown that such view can inherit important structural elements from game theoretic models of bargaining – in particular, for the case where the protagonists have overlapping negotiation regions, we have emphasized a parallel to the Nash solution in cooperative game theory. When acceptable solutions regions of the protagonists are disjoint, we have presented several types of processes: changes in the salience of dimensions, dimensional projections and metaphorical space transformations. None of the latter processes are motivated by normative or rationality considerations, but presented as argumentation tools that we believe are used in actual situations of conceptual disagreement.

Acknowledgements

References

Andersson, T. (1994). *Conceptual Polemics: Dialectic Studies of Concept Formation*. Lund: Lund University Cognitive Studies 27.

Barker, C. (2002). The dynamics of vagueness. *Linguistics and Philosophy*, *25*(1), 1–36.

Barker, C. (2013). Negotiating taste. *Inquiry*, *56*(2-3), 240–257.

Bates, E. (1976). *Language and Context: The Acquisition of Pragmatics*. New York: Academic Press.

Brennan, S. (1996). Lexical entrainment in spontaneous dialog. In *Proceedings, 1996 International Symposium on Spoken Dialogue, ISSD-96. Philadelphia, PA* (pp. 41–44).

Brennan, S. E., & Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*(6), 1482–1493.

Brinck, I. (2004). The pragmatics of imperative and declarative pointing. *Cognitive Science Quarterly*, *3*(4), 429–446.

Clark, H. H. (1996). *Using Language*. Cambridge: Cambridge University Press.

Clark, H. H., & Schaefer, E. F. (1989). Contributing to discourse. *Cognitive Science*, *13*(2), 259–294.

Egré P. (1913) "What's in a planet?" In M. Aloni, M. Franke and F. Roelofsen eds., *The Dynamic, Inquisitive and Visionary Life of phi, ?phi and Diamond-phi, A Festtschift for J. Groenendijk, M. Stokhof and F. Veltman*, ILLC, 2013, pp. 74-82.

Furnas, G. W., Landauer, T. K., Gomez, L. M., & Dumais, S. T. (1987). The vocabulary problem in human-system communication. *Communications of the ACM*, *30*(11), 964–971.

Gärdenfors, P. (2000). *Conceptual Spaces: The Geometry of Thought*. Cambridge, MA: MIT Press.

Gärdenfors, P. (2014). *The Geometry of Meaning: Semantics Based on Conceptual spaces*. Cambridge, MA: MIT Press.

Gärdenfors, P. and Warglien, M. (2013). The development of semantic space for pointing and verbal communication. In J. Hudson, U. Magnusson, & C. Paradis (Eds*.), Conceptual Spaces and the Construal of Spatial Meaning: Empirical Evidence from Human Communication* (pp. 29-42). Cambridge: Cambridge University Press.

Garrod, S., & Anderson, A. (1987). Saying what you mean in dialogue: A study in conceptual and semantic co-ordination. *Cognition*, *27*(2), 181–218.

Grice, H. P. (1989). *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.

Harsanyi, J. C. (1956). Approaches to the bargaining problem before and after the theory of games: A critical discussion of Zeuthen's, Hicks', and Nash's theories. *Econometrica*, *24*(2), 144–157.

Larson, R. K., & Ludlow, P. (1993). Interpreted logical forms. *Synthese*, *95*(3), 305–355.

LiCalzi, M. and Maagli, N. (2013). Bargaining over a common concptual space, manuscript.

Ludlow (to appear*). The Dynamic Lexicon*, manuscript.

McNamara, R. S., Biersteker, R. S. M., James Blight, Robert K. Brigham, Thomas J., Blight, J., Brigham, R. K., Biersteker, T. J., & Schandler, C. H. (2007). *Argument Without End: In Search of Answers to the Vietnam Tragedy*. New York, NY: Public Affairs.

Nenkova, A., Gravano, A., & Hirschberg, J. (2008). High frequency word entrainment in spoken dialogue. In *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics on Human Language Technologies: Short Papers* (pp. 169–172). Stroudsburg, PA, USA: Association for Computational Linguistics.

Ortony, A., Vondruska, R. J., Foss, M. A., & Jones, L. E. (1985). Salience, similes, and the asymmetry of similarity. *Journal of Memory and Language*, *24*(5), 569–594.

Osborne M. and A. Rubinstein (1994), *A Course in Game Theory*, Cambridge, MA: MIT Press

R. Parikh. Vagueness and utility: the semantics of common nouns. *Linguistics and Philosophy*, 17: 521-535.

Perelman, C., & Olbrechts-Tyteca, L. (1969). *The New Rhetoric: A Treatise on Argumentation*. Notre Dame, IN: University of Notre Dame Press.

Pickering, M. J., & Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, *27*(02), 169-190

Pinker, S., Nowak, M. A., & Lee, J. J. (2008). The logic of indirect speech. *Proceedings of the National Academy of Sciences*, *105*(3), 833–838.

Reitter, D., & More, J. D. (2007). Predicting success in dialogue. *Proceedings of the 45th Annual Meeting of the Association for Computatinal Linguistics* (pp. 808-815), Prague: Association for Computational Linguistics.

Rocci (2009). Maneuvering with voices. In Van Eemeren, F. H. (2009). *Examining Argumentation in Context: Fifteen Studies on Strategic Maneuvering*. John Benjamins Publishing.

Ryan, D. (2004). The pragmatic theory of meaning: Negotiation by stealth. *Language Sciences*, *26*(2), 217–229.

Schelling, T. C. (1960). *The Strategy of Conflict*. Cambridge, Massachusetts: Harvard University.

Selten, R & Warglien, M. (2007). The emergence of simple languages in an experiemional coordination game. *Proceedings of the National Academy of Sciences*, *104*(18), 7361–7366.

Stalnaker, R. (1999). *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford: Oxford University Press.

Taylor, S. E., J. Crocker, S.T., Fisle, M. Sprinzen, & J.D. Winkler (1979). The generalizability of salience effects. *Journal of Personality and Social Psychology*, *37*(3), 357–368.

Thomson, W. (1994). Cooperative models of bargaining. In *Handbook of Game Theory with Economic Applications* (Vol. 2, pp. 1237–1284). Amsterdam: Elsevier.

Van Benthem, J. (2008). "Games that make sense": logic, language, and multi-agent interaction. In *New Perspectives on Games and Interaction* (pp. 197–210). Amsterdam: Amsterdam University Press.

Van Eemeren, F. H. & Grootendorst, R. (2004). *A Systematic Theory of Argumentation: The Pragma-Dialectical Approach*. Cambridge: Cambridge University Press.

Van Eemeren, F. H.(2010) *Strategic Manoevering in Argumentative Discourse*. Amsterdam: John Benjamins Publishing Company.

Von Ahn, L. (2006). Games with a purpose. *Computer*, *39*(6), 92–94.

Warglien, M., & Gärdenfors, P. (2013). Semantics, conceptual spaces, and the meeting of minds. *Synthese*, *190*(12), 2165–2193.

Waismann, F. (1968). "Verifiability", in A.G..N. Flew (ed.), *Logic and Language*, pp. 117-144, Oxford: Blackwell.

Wenger, E. (1998). *Communities of Practice: Learning, Meaning, and Identity*. Cambridge: Cambridge University Press.

Zeuthen, F., Zeuthen, F. E., & Wiggs, K. I. (1930). *Problems of Monopoly and Economic Warfare*. London: Routledge & Sons.