Proceedings of the 18th Yale Workshop on Adaptive and Learning Systems
Yale University, New Haven, 2017

Adaptive Procedures in Conversational Interactions

Matthias J. Feiler◇

***Abstract*—In this paper we describe how participants in a conversation understand each other by recognizing in each others' statements episodes of their own experience. We observe that due to the underspecification of propositions the congruence with a participant's own views is overestimated. We demonstrate that this bias enables speakers to come to a mutual agreement on a blended concept which combines elements – but lives outside of the speakers' individual information sets.**

***Index Terms*—mutual adaptation; cooperative cognition; blending**

I. Introduction

When people communicate, they rely on conventions in order to understand and produce meaning. Meaning is constructed in the mind of the listener using language as an input from which conceptual representations are formed. These linguistic inputs typically under–specify the concepts intended by the speaker and rely on the listener's ability to contribute the context needed to make a correct inference. In rational interaction models the speaker and listener apply (and expect) a common logic, or *cooperative principle* [1] to organize their speech acts.

The principle has been spelled out into four conversational maxims, the maxim of quality (truthfulness), quantity (informativeness), relevance and manner (conciseness). Mutual agreement on the maxims allow the speaker and listener to enrich an utterance by so–called *implicatures* which suggest an extension or modification of meaning beyond the literal interpretation, such as in *S1:* "Will he come?" *S2:* "His car broke down." which is decodable by *S1* into "He won't." by assuming that *S2* did not choose the answer if it was irrelevant. Also, *S2* supposes that *S1* has the background information that if cars brake down, people frequently do not manage to keep appointments. This is referred to as the *common ground* [2], [3]. The interactive alignment model (IAM) [4] emphasizes the importance of tacit coordination and *implicit* common ground. According to the model, grounding occurs automatically and the speakers' particular choices (i.e. which information to foreground) lead to an alignment of their (mental) representations.

Following a long tradition [5], conversation analysts study the way an interaction order [6] is established in practice, in particular how people take turns at talk, how they deal with overlaps and interruptions and how the sequence of utterances (and more general [speech] actions) is organized. Conversation analysis argues that the "...meaning of an action is heavily shaped by the sequence of actions from which it emerges, and that the social context is dynamically created [...] through the sequential organization of interaction", see [7], p.223. Any statement has to signal understanding of the preceding statements and prepare the floor for the next in order to establish coherence. This means that "each sentence [...] must contain some direct or indirect indication as to how it fits into the stream of talk", see [8], p.119. Two minds have to collaborate in order to "make progress" on the subject of their discussion. In the IAM this process occurs with a minimal amount of modeling what others know.

The formulation of the IAM provoked a prolific discussion on the role of alignment and the proposed underlying priming mechanism that creates it. For the purpose of this paper we wish to retain the following idea: the understanding of an utterance depends on the relation of the proposition with the listener's own representation of the situation (i.e. his or her situation model [9]). If the models are aligned the information is more easily accessible for the listener. However, as we shall demonstrate, dialogue can still be successful if the speaker and listener maintain different underlying representations. In fact, these differences might be a characteristic of creative discourse, which gives rise to a new representation for both participants.

Conversations are an example of a more general class of interactions in which participants operate symbiotically to achieve a greater goal at the group–level. In contrast to consensus problems [10], participants must not simply follow others but constantly decide whether to adopt or oppose a given view. They coordinate in a way to avoid complete alignment as this would compromise the "give–and–take" nature of their symbiotic arrangement. As a consequence, the convergence mechanism in this type of interaction is more intricate than the one obtained when applying rules such as averaging–among–neighbors [11]. When agents have to agree to disagree, it is not clear what the final arrangement will be. In [12], a problem of distributing agents over a set of different unknown values led to a complex convergence pattern in which the final assignment of agents is unclear until a neighborhood of the values is reached by all agents simultaneously. More recently, instability has been discovered in systems of models that attempt to stabilize each other through adaptation [13], [14]. In summary, new adaptive rules have to be devised which ensure that agents *selectively* exchange information in such a way as to form a meaningful (in fact, conversation–like) sequence. We will refer to this problem as mutual adaptation.

II. From perception to concept formation in dialogue

In this paper, we view conversations from a social influence [15] perspective. We argue that the pressure towards alignment (if any) originates in the participants' desire to *belong* to the

◇E-mail: matthias.feiler@uzh.ch

conversation process and conform (to some extent) with the "rules" of that process. The very fact of initiating an exchange of ideas deemed relevant by the participants establishes an in–group of people sharing those ideas. Participants are ready to adjust their own views in order to even out possible differences with in–group members.

The effect may be thought of as a variant of the foot–in–the–door technique [16] in that a lock–in occurs after the first successful encounter. Utterances that are perceived as relevant and congruent with one's own views are thought to belong to "like–minded" speakers. An anchoring effect [17] occurs because the label *like–minded* belongs to the *speaker* and not to the utterance produced. The initial perception of likeness creates a bond between speaker and listener that both parties seek to retain even as further details (coming up as the conversation develops) may reveal a disagreement in beliefs.

Once speaker and listener have established a relation (of like–mindedness) expectations on the content and attitude expressed in future statements are formed. If these expectations are not fulfilled, the participants experience cognitive dissonance, a discomfort caused by the inconsistency of the expected vs. realized continuation of the dialogue. At this point one might ask what strategies are available to resolve this tension? We argue that the interlocutors will adapt some positions if they are accessible [18] from their respective backgrounds.

In order to clarify the idea let us identify linguistic units with points in the two-dimensional plane as in [19]. The units are assembled by the beholder to form a holistic picture corresponding to the conceptual representation of the sensory inputs received. In other words, the mind relates to inputs in terms of an aggregate structure which becomes meaningful through its global *shape*. This is the viewpoint developed by Gestalt psychologists who also established the basic principles governing the assembly of shapes [20]. In this paper we assume that the conversation partners differ in the way they "connect the dots" received in an speech event. For simplicity, we assume that speaker 1 is a SQUARE–thinker, meaning that he or she tends to mentally group individual points into the *shape* of a square and speaker 2 is a CIRCLE–thinker. We will refer to these two primitive shapes as the background frames, a kind of knowledge base used in conversations in order to *access* the conceptual propositions presented by the other speaker.

Coming back to the question of how to resolve dissonant (i.e. unexpected) propositions we assume that the conversation partners activate their respective shape repertoire in order to try and accommodate the new situation created by an off-key statement. The circle thinker will adjust his or her view by "putting a circle around incongruent dots" thereby finding a new composition of primitive shapes that restores the level of agreement initially observed with the other speaker. Put differently, in a circle–world conceptual representations are adapted by adding or subtracting circles (of varying size) to existing ones. The same holds true for the square thinker (speaker 1). If both speakers adapt their positions during the course of the conversation to find greater agreement and at the same time stick to their basic building blocks for mental representations (actually speakers are constrained by them) a mixing of concepts occurs, in the sense that squares and circles cross paths at points where a disagreement is being resolved. We regard this situation to be naturally conducive to conceptual blending [21] as participants are in essence trying to fit a square into a circle (and vice versa).

The interesting point to note is that both speakers contribute to the formation of the blend without giving up their respective mental frames. This construction is possible because of a perceptual similarity experienced by the speakers. As in [22], perceptions are *tuned* (i.e. adapted) to meet the demands of a successful conversation. At the same time, a new, integrated mental space emerges at the discourse level, which –at some point– will be recognized by the participants as an independent structure having a meaning of its own. In our view, this point corresponds to the *creative moment* witnessed by the conversation partners (often simultaneously).

III. The Conversation Model

In our model with two participants, turn–taking is organized as a simple alternating sequence, i.e. a participant's statement is followed by a response of the other participant which again triggers a statement etc. The conversation begins with a statement that triggers the interest of a listener. By that we mean that the listener perceives the statement as a fragment of a larger, "bigger picture" to which he or she relates. After the first interaction both speaker and listener have an idea of the scope of the conversation. The scope constrains the set of eligible linguistic expressions used for describing concepts that lie in the domain of interest. We assume that both partners have access to the same (finite) set of expressions and state our objective as follows:

Statement of the objective: Let S be a finite set of linguistic units available to both speakers. The speakers differ in their assessment which of the available units are constitutive of the concept under discussion. They access subsets of S in order to build mental representations of the concept. The sets of subsets used by the two speakers are different. The objective is to arrive at a mutual agreement on the definition of the concept by combining subsets of S.

A. Description

As indicated above, we identify concepts with two–dimensional geometric objects which are composed of simple *building blocks* which, in turn, are constituted from *points* corresponding to linguistic units. We assume that the first speaker represents world concepts using SQUARES of all sizes and the second speaker uses CIRCLES of all sizes. While these shapes correspond to the mental representations (constructed from building blocks), only points are disclosed during a conversation. For example, if the second speaker describes a circle he or she will disclose a number of points along the circumference. Any fellow CIRCLE–thinker would see a circle but our assumed conversation partner will see elements of an "outer" SQUARE in which the circle is embedded. This of course simply means that utterances do not fully specify concepts but rely on the ability of listeners to "fill the gaps". Our model proposes three steps towards mutual understanding.

Step 1: Comprehension. If the SQUARE in the above example is confronted with points obtained from a circle he or she will continue to believe in the square until substantial counter–evidence (in a sense defined more precisely below) is presented. As long as the listener erroneously believes in the wrong shape he over–estimates the agreement with the speaker. Once too many inconsistent statements are received, the SQUARE adopts some of the opposite views.

We assume that the recognition of differences proceeds in the same way as the perception of agreement. As long as the differing positions are scattered randomly in the plane they will be seen as outliers and overlooked. When differences accumulate around certain points they will be recognized as larger units. Listeners "package" the affected points thereby assigning a *structure* to the area of disagreement. This is the point where the evidence against an initial representation (e.g. a simple square) is "substantial". Acting on the Gestalt principle of closure [23] a shape is assigned to the area making it accessible to inclusion or exclusion from the listener's mental model. In our geometric analogy the listener will find the smallest square which covers all unequal positions and add or subtract it from the existing model. The same holds for representations consisting of circles.

Step 2: Adaptation. The central element of our model is that participants seek to retain their initial agreement level even as evidence of disagreement becomes clear in course of the conversation. They change some of their positions and adopt the partner's stance in an attempt to keep the conversation alive. The partner mirrors this behavior and adjusts some of his or her positions following a tit for tat [24] logic. The participants weigh the cost of giving up positions against the benefit of continuing the conversation.

As described in step 1, the adjustments occur from a position of strength in which the conversation partners focus on the confirmations received by the other speaker. Instead of reacting to every single contradicting input listeners simply ignore them until a sufficiently large number accumulates which will be treated as a unit. From an information processing perspective this dramatically reduces the number of decisions necessary in the negotiation of understanding. Adjustments are made in a step–wise fashion whenever the current model cannot be sustained. The cost of adjusting is further reduced as speakers only take "native steps", i.e. they add or subtract shapes belonging to their natural vocabulary of mental building blocks. This means that they do not fully adjust the way they think about a proposed concept but they merely *internalize* those aspects that can be expressed by their own native language. Even as positions are adjusted the speakers continue to operate under the hypothesis of "being understood" in their respective contexts (squares or circles). In summary, the adaptation of views proceeds *subtractively* starting from an initially (coarse–grained) perceived match of positions.

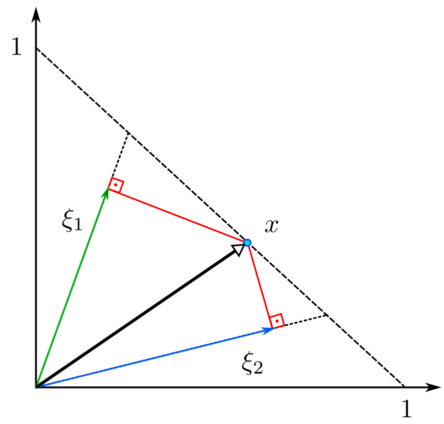

Fig. 1. Geometry of the representations ξ_1 and ξ_2 of the blended concept x

Part 3: Production. In our model, the production of utterances corresponds to sampling from a participant's current mental representation of the concept under discussion. Points are drawn from *interesting* areas of the shape that is under construction in the minds of the interlocutors. If positions are adapted the corresponding speaker will signal his "change of mind" by sampling primarily from areas containing updated positions. Signaling is critical as it precludes the simultaneous adjustment of positions by both parties which would result in an oscillation between opposite stances.

Our turn–taking schedule ensures that the role of speaker and listener is always well-defined. If listeners incorporate positions of the speaker and signal this during the next turn, the partners are able to make progress on the subject of their conversation. As a result, the areas in which adjustments are necessary will eventually become smaller as the scope of the conversation is being exhausted. An equilibrium is reached if a minimum block size for representing the agreement is reached by both participants. In other words, no further areas of disagreement may be covered (and, hence, adapted) by available building blocks. As no further areas of disagreement are discernable, the participants have reached the maximal agreement possible within the (finite) resolution of their combined cognitive abilities.

B. Formalization

We assume that our conversation takes place in some n–dimensional space where each dimension corresponds to a linguistic unit. Concepts are represented as vectors $x \in \mathbb{R}^n$ whose elements correspond to the weight (importance) of every linguistic unit in the definition of the concept. Speaker 1 and 2 "live" in subspaces $S_1, S_2 \subset \mathbb{R}^n$ of substantially lower dimensions. $P_1 x \in S_1$ is the projection of a point $x \in \mathbb{R}^n$ on the subspace S_1 corresponding to the aspects of x that speaker 1 is able to comprehend. Also, the linguistic output of speaker 1 will be constrained to S_1. If S_1 and S_2 are orthogonal, speaker 1 and 2 are unable to communicate as $P_2 P_1 x = 0$ for any description $P_1 x$ of x by speaker 1.

In general, we assume that $P_i P_j x \neq 0$ for $i \neq j$ and $x \neq 0$. Given two representations $\xi_1 \in S_1$ and $\xi_2 \in S_2$ there exists (at least one) $x \in \mathbb{R}^n$ such that $P_1 x = \xi_1$ and $P_2 x = \xi_2$. We will refer to x as a blend achieved through conversation, or ***conversational blend*** if the following conditions hold:

Definition: x is a conversational blend if

- $\|x\|_1 = 1$, and
- $P_i x = \xi_i$ for $\xi_i \in S_i \quad i = 1, 2$.

The first condition states that the weights are distributed over x such that the sum of weights equals 1. This may be realized by counting the occurrence of any specific linguistic unit relative to the total number of units employed by a speaker. The second condition relates x to its mental representations ξ_1, ξ_2 by speaker 1 and 2.

At start of the conversation, both speakers (erroneously) believe that the topic under discussion is exactly what they think it is, i.e. $x_i = \xi_i$ for $i = 1$, 2. It is clear that they cannot come to a mutual agreement since $P_i x_j \neq \xi_i$, $j \neq i$ unless the two subspaces coincide. The speakers gradually realize that the topic x_i and their idea (mental representation) of the topic ξ_i are different. They will search for a new pair $(x_i,\, \xi_i)$ until the above blending conditions are met. The interesting point to note is that the equilibrium $x^* = x_1^* = x_2^*$ will necessarily lie *outside* (at least one of) the subspaces giving rise to the defining property of x as a blend. Also, at least one of the vectors ξ_i will have a norm less than one. This means that ξ_i no longer defines a distribution over linguistic units. Reducing the length of ξ_i is equivalent to acknowledging that x cannot be modeled as a distribution over a speaker's own linguistic repertoire. Some aspects lie beyond his or her mental grasp. It is precisely this concession of incompleteness that enables speakers to arrive at an agreement on the blended concept. At the same time, ξ_i remains in the subspace S_i since linguistic units are added or subtracted in packages of native building blocks.

The equilibrium condition is depicted in figure (1) for the case of $n = 2$ linguistic units. In practice, n will be very large and every utterance corresponds to a sample drawn from the low–dimensional subspace $S_i \subset \mathbb{R}^n$ of building blocks (sets of linguistic units) accessible to a speaker. As more and more samples are disclosed speakers gradually become aware of the location of concept x_1 relative to x_2 giving rise to the adaptation of ξ_i, $i = 1, 2$.

Comment: In the formal definition above we assumed that the "scope" of the conversation is known to both speakers as the set of linguistic units available for describing relevant concepts. This enables us to define the norm relative to the same set linguistic units and, in particular, to display both speakers in the same figure (1). In practice, every speaker will have his or her own (finite) collection of relevant linguistic units. This does not defy the above arguments but simply means that some dimensions of $\mathbb{R}^n$ are used by one but not the other speaker thereby constraining the corresponding elements in ξ to zero.

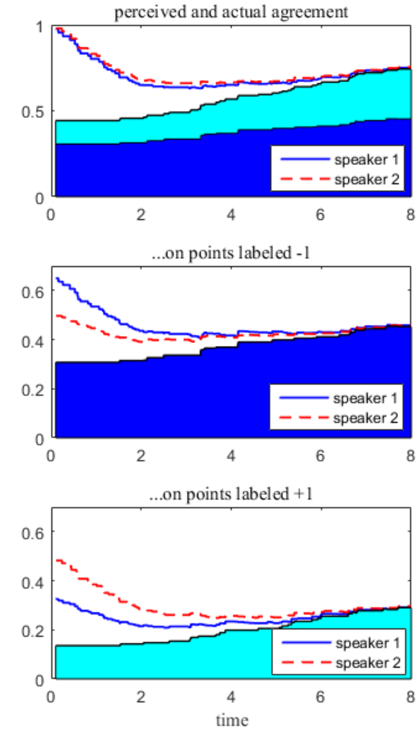

Fig. 2. Evolution of the agreement among conversation partners as measured by the number of congruent grid labels divided by the total number of grid points. The shaded areas correspond to the *ground truth* while the lines report the agreement level as perceived by the participants.

IV. Simulation

We use computer simulations to illustrate how conversational interactions occur in our model. The aim is to demonstrate that blends occur naturally if the assumptions of our model are satisfied. Real–life examples include board–room discussions where blends may be a way of settling differences in consensus decisions. Another source of empirical data (to be used in future studies) may be found in online (weblog) conversations where the cooperative principle is often neglected due to the absence of the coordination mechanisms available in classic conversations. If disconnected messages are simply juxtaposed the reader will either have to filter out or blend parts of the messages in order to resolve the inconsistency.

Speech events. In our simulation, utterances correspond to points drawn randomly (with replacement) from a grid that is assumed to be common to both speakers. The points are labeled ($+1$ and -1) indicating whether or not they are part of the two–dimensional object under discussion. The discussants differ in their respective labeling of grid points. During the course of the conversation some of the labels are updated in an effort to sustain mutual understanding. Speech events are governed by a stochastic Bernoulli process in which the success probabilities alternate from a high to a low value according to a simple turn–taking schedule as in [25]. The high value passes from one participant to the other such that

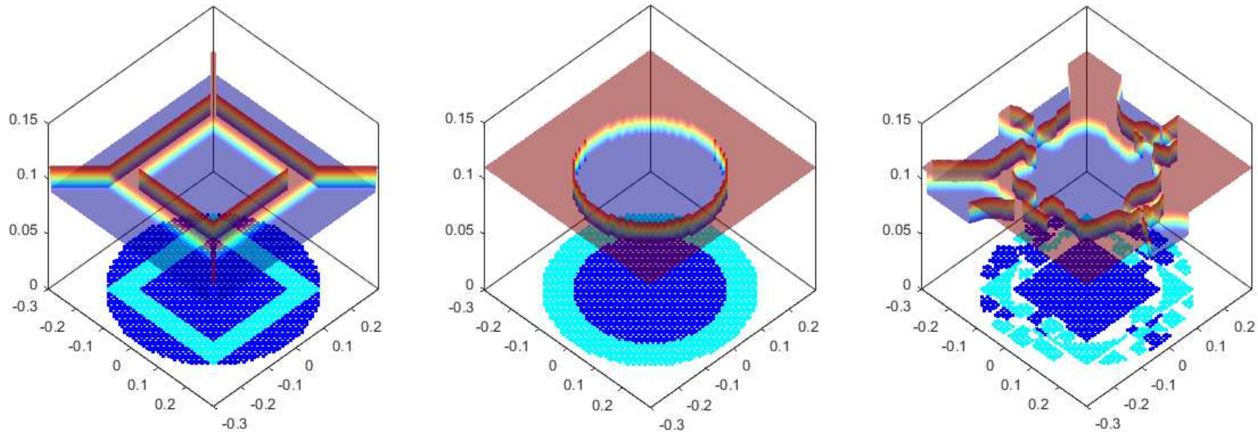

Fig. 3. (From left to right:) Input 1 (square frame), Input 2 (ring) and blended space (mix). Points belonging to the object are labeled $+1$ and marked as bright dots and non–members are labeled -1 corresponding to dark dots in the 2-D projection below the object. In the blend, an agreement on members and non–members has been achieved among the conversation partners.

the roles of speaker and listener are well–defined at every instant of time.

Classification. Our model participants recognize and process points in packages of larger units corresponding to the frames used in the mental representation of incoming information. In our computer implementation each speaker is endowed with a maximum margin classifier [26] which employs the respective native building blocks as a decision boundary. The blocks are combined to obtain a map of the agreement and disagreements (on labels). Whenever a majority of points within a block contradicts a speakers's own positions he or she will flip *all* positions within the block. The majority rule is critical as it guarantees that the total number of agreed points in the conversation does not decrease.

At every turn, speakers attempt to resolve differences. We assume that the block size (diameter of the square or circle) may be reduced until the majority condition is met. Also, after converting positions within an appropriately sized block, participants disclose their new stance by using the points with updated labels as part of their next utterance. It is clear that these assumptions may not always hold in practice. On the other hand, the algorithm mimics a natural tendency of conversations to develop focus (areas) over time. If this focus is on differing positions the pressure for one of the sides to adjust increases. The (uninteresting) case excluded from our analysis is that initially compatible conversations may "get stuck" at details if none of the parties moves.

Conversational flow. The scope of the conversation is defined by the first batch of samples drawn by a speaker from a bounded region of an (infinite) two–dimensional grid. A second speaker engages if the samples lie within his or her region of interest. At this point two objects appear in the minds of the speakers and both think their own representation is what the other speaker means. Figure 3 (left) displays two such objects (a ring and a square frame) present at the onset of the discussion. With every utterance produced by the other speaker his or her *actual* positions become evident. When assessing the agreement level, our participants measure *by subtraction* i.e. they retain their own representation as a reference and subtract incongruent statements received from the other. Figure (2) displays the bias towards higher perceived agreement levels resulting from this operation. The agreement level in our simulation is defined as the number of matching labels over the total number of grid points within the scope. We report the agreement on both member (positive labels) and non–member points (negative labels).

Whenever a disagreement accumulates at a certain point it may be captured by one of the speakers who will flip his or her positions in an attempt to restore the original perceived agreement level. This gives rise to the "hockey–stick" recovery observed in the evolution of agreement in figure (2). The process of resolving differences continues even after all positions have been disclosed. The participants' target agreement level is *anchored* at the initial perception of (almost) perfect agreement.

Figure (4) displays the resulting compromise after many encounters. Participants have in– and excluded points to and from their original member sets in packages corresponding to circles and squares. We assume a minimum package size which means that the member sets of the two speakers cannot be fully aligned. An equilibrium is reached once this minimum size is attained, i.e. packages cannot be further reduced to ensure the majority condition (as defined above). Due to the stochastic nature of the sampling process the equilibrium outcome, i.e. the final shape achieved through blending, is not known a priori. Figure (3) displays one such outcome.

V. Conclusion

We argue that adaptive procedures towards mutual understanding in conversations are driven by the participants'

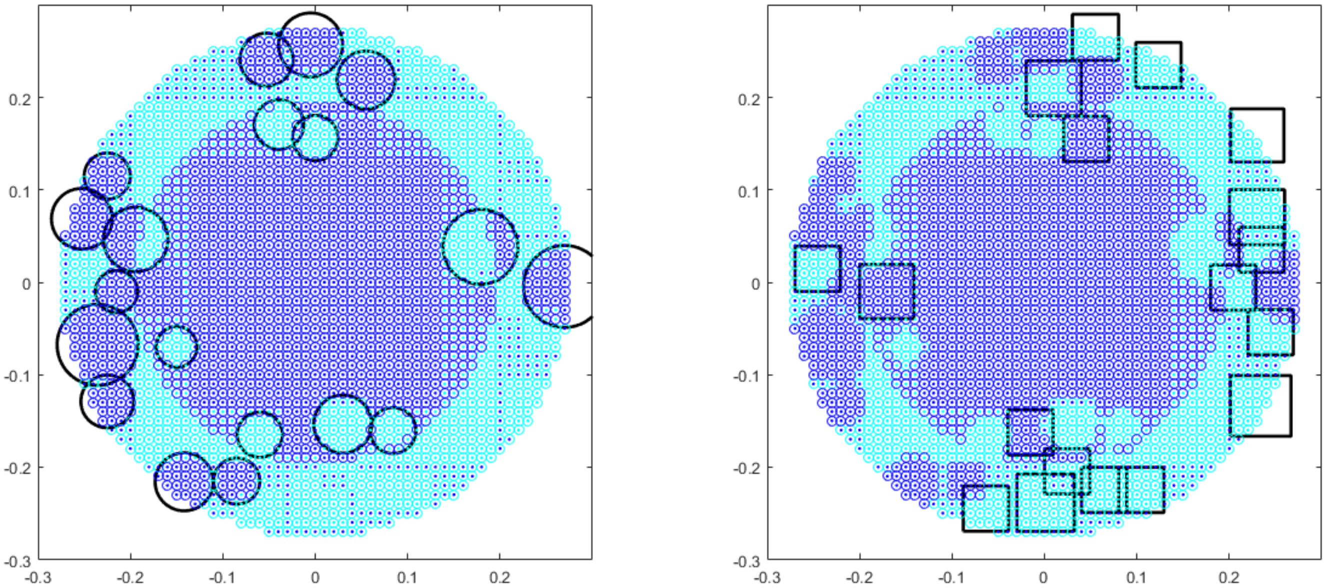

Fig. 4. Mental representation of the blended object by speaker 1 (left) and 2 (right). Every circle (square) corresponds to an area of previous disagreement which has been identified and updated asynchronously by the conversation partners.

effort to restore the agreement level perceived at the onset of the conversation. We demonstrate how conceptual blends arise when the adaptation of positions is constrained by the participants' "native" ways of mentally representing concepts: a SQUARE will complete a set of points differently from a CIRCLE–thinker. A stable equilibrium corresponds to a shape in which all differences that can be captured by native patterns have been removed by the interlocutors.

In contrast to other learning schemes e.g. the Bayesian update rule we assume that speakers do not "switch" to the other representation (i.e. from square to circle or vice versa) even as evidence supporting the other shape accumulates simply because they cannot access the other representation using native building blocks.

REFERENCES

[1] H. P. Grice, "Logic and conversation," *1975*, pp. 41–58, 1975.

[2] R. Stalnaker, "Common ground," *Linguistics and Philosophy*, vol. 25, no. 5-6, pp. 701–721, 2002.

[3] H. H. Clark and C. R. Marshall, "Definite reference and mutual knowledge," *Psycholinguistics: critical concepts in psychology*, vol. 414, 2002.

[4] M. J. Pickering and S. Garrod, "Toward a mechanistic psychology of dialogue," *Behavioral and brain sciences*, vol. 27, no. 02, pp. 169–190, 2004.

[5] E. A. Schegloff, *Sequence organization in interaction: Volume 1: A primer in conversation analysis*. Cambridge University Press, 2007, vol. 1.

[6] E. Goffman, *Forms of talk*. University of Pennsylvania Press, 1981.

[7] J. Heritage, "Conversation analysis and institutional talk," *Handbook of language and social interaction*, pp. 103–147, 2005.

[8] J. J. Gumperz, "Mutual inferencing in conversation," *Mutualities in dialogue*, pp. 101–123, 1995.

[9] R. A. Zwaan and G. A. Radvansky, "Situation models in language comprehension and memory." *Psychological bulletin*, vol. 123, no. 2, p. 162, 1998.

[10] A. Jadbabaie, J. Lin, and A. S. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Transactions on automatic control*, vol. 48, no. 6, pp. 988–1001, 2003.

[11] D. P. Bertsekas and J. N. Tsitsiklis, "Comments on coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Transactions on Automatic Control*, vol. 52, no. 5, pp. 968–969, 2007.

[12] M. J. Feiler and K. S. Narendra, "Simultaneous identification and control of time-varying systems," in *Decision and Control, 2006 45th IEEE Conference on*. IEEE, 2006, pp. 1093–1098.

[13] K. S. Narendra and P. Harshangi, "Unstable systems stabilizing each other through adaptation," in *American Control Conference (ACC), 2014*. IEEE, 2014, pp. 7–12.

[14] ——, "Unstable systems stabilizing each other through adaptation-part ii," in *American Control Conference (ACC), 2015*. IEEE, 2015, pp. 856–861.

[15] J. C. Turner, *Social influence.* Thomson Brooks/Cole Publishing Co, 1991.

[16] J. L. Freedman and S. C. Fraser, "Compliance without pressure: the foot-in-the-door technique." *Journal of personality and social psychology*, vol. 4, no. 2, p. 195, 1966.

[17] A. Tversky and D. Kahneman, "Judgment under uncertainty: Heuristics and biases," in *Utility, probability, and human decision making*. Springer, 1975, pp. 141–162.

[18] J. S. Bruner, "On perceptual readiness." *Psychological review*, vol. 64, no. 2, p. 123, 1957.

[19] G. Radden, K.-M. Köpcke, T. Berg, and P. Siemund, *Aspects of meaning construction*. John Benjamins Publishing, 2007.

[20] K. Koffka, *Principles of Gestalt psychology*. Routledge, 2013, vol. 44.

[21] G. Fauconnier, "Conceptual blending and analogy," *The analogical mind: Perspectives from cognitive science*, pp. 255–286, 2001.

[22] R. L. Goldstone and L. W. Barsalou, "Reuniting perception and conception," *Cognition*, vol. 65, no. 2, pp. 231–262, 1998.

[23] M. Wertheimer, "Laws of organization in perceptual forms." 1938.

[24] R. Axelrod *et al.*, "The evolution of strategies in the iterated prisoners dilemma," *The dynamics of norms*, pp. 1–16, 1987.

[25] M. Feiler, "An agent model of conversation," *Proc. of the 17th Yale Workshop on adaptive and learning systems*, 2015.

[26] R. E. Schapire, Y. Freund, P. Bartlett, W. S. Lee *et al.*, "Boosting the margin: A new explanation for the effectiveness of voting methods," *The annals of statistics*, vol. 26, no. 5, pp. 1651–1686, 1998.