Relevance for Dialogue

Jonathan Ginzburg
King's College, London
London, UK
jonathan.ginzburg@kcl.ac.uk

Abstract

Relevance in the sense of *conversational coherence* is the most fundamental notion for research on dialogue. It is the cornerstone of theories of dialogue in the same way that *grammaticality* is to syntax. In this paper I restrict attention to relevance relating a query to a possible (felicitous) response. Still, even restricted to this domain, attempts at a comprehensive characterization of relevance, difficult as they undoubtedly are, are few and far between. Indeed, most existing accounts are intrinsically restricted in their ability to scale up. I offer a number of arguments for the need for a notion of relevance internalized in some way within the theory of meaning: relevance seems to underpin certain types of clarification questions and *lack of conversational relevance* seems to underpin the inference that one does not wish to address a prior utterance. I sketch an account of relevance within the dialogue theory KoS underpinned by Type Theory with Records.

1 Introduction

Relevance in the sense of *conversational coherence* is the most fundamental notion for research on dialogue. It is the cornerstone of theories of dialogue in the same way that *grammaticality* is to syntax. Indeed (Turing, 1950) proposed that the ability to evince relevance in approximately this sense could be a plausible test for intelligence. In what follows, I restrict attention to relevance relating a query to a possible (felicitous) response.[1] This is, in part, due to obvious considerations of space, but also because this is a domain where considerable work has been done on one component of the problem. Still, even restricted to this domain, attempts at a comprehensive characterization of relevance, difficult as they undoubtedly are, are few and far between. Indeed, as I will explain below, most existing accounts are intrinsically restricted in their ability to scale up.

Beyond this, one issue to consider is whether there really is a need for a single notion of relevance—whose restriction to query moves we discuss here, internalized in some way within the theory of meaning.[2] There is at least one substantive argument for internalizing relevance, as well as some methodological motivation. The substantive argument is that a unitary notion of conversational relevance seems to underpin certain types of clarification questions—ones that arise when the coherence of an utterance seems unclear, as in ((1)a).[3] Similarly, as Grice famously pointed out, *lack of conversational relevance* seems to underpin the inference that one does not wish to address a prior utterance, as in ((1)b):

(1) a. Marjorie: Don't touch that cos she hasn't had it yet.

[1] For a more detailed account see (Ginzburg, 2011).

[2] I use 'theory of meaning' to avoid boring territorial disputes between semantics and pragmatics. I attempt, nonetheless, to be reasonably explicit as to whether components of the theory of relevance refer to public context or to agent–internal parameters, which is essentially how I view the distinction.

[3] An empirical caveat is in order here. 'What do you mean' is clearly NOT a purpose built CR for querying the relevance of an utterance. In practice, the vast majority of 'what do you mean' CRs, at least in the BNC, seem to be about literal content, NOT about coherence:

(i) Anon 6: No, there's nobody here much Richard: What do you mean there's nobody here, it's packed.

(ii) Cassie: You did get off with him? Catherine: Twice, but it was totally non-existent kissing so Cassie: What do you mean? Catherine: I was sort of falling asleep.

Pretheoretically this is perhaps not surprising, given that (perceived) complete lack of coherence is rare; whereas indeterminacy of content is a consequence of lexical and phrasal context dependence, but a detailed explanation is surely an important desideratum.

Aspects of Semantics and Pragmatics of Dialogue. SemDial 2010, 14th Workshop on the Semantics and Pragmatics of Dialogue

Dorothy: Does she eat anything? Marjorie: What do you mean? (British National Corpus (BNC))

b. Dr. Grimesby Roylott: My stepdaughter has been here. I have traced her. What has she been saying to you?
Sherlock Holmes: It is a little cold for the time of the year.
Dr. Grimesby Roylott: What has she been saying to you?
Sherlock Holmes: But I have heard that the crocuses promise well. ('The Speckled Band', Sir Arthur Conan Doyle, *The Adventures of Sherlock Holmes*, John Murray, London.)

The methodological argument is that characterizing relevance pushes theories of dialogue to be concrete, forcing them to be precise about the range of propositions they characterize as answers and to offer sources of relevance to utterances whose relevance as an answer they do not underpin. It also enables one to operationalize the notion of relevance for use in corpus studies and for other computational work.

In this paper I sketch an account of relevance within the dialogue theory KoS (Ginzburg, 1994; Ginzburg and Cooper, 2004; Larsson, 2002; Purver, 2006; Fernández, 2006; Ginzburg, 2011; Ginzburg and Fernández, 2010). The basic intuition is that relevance of an utterance relative to an agent's information state amounts to the possibility of integrating the utterance into the information state; though as we will see this basic intuition needs to be refined to deal with cases like ((1)b). I start by offering a more or less theory neutral characterization of relevance, suggesting the need to encompass (in approximate order of theoretical difficulty)

- *q*(uestion)-specificity—this includes both answerhood and some sort of dependence or entailment relation between questions,
- metadiscursive relevance (a notion that underwrites utterances like "I don't know" and 'I don't want to talk about this.")
- genre-based relevance, the latter much studied in AI work on dialogue
- metacommunicative relevance, a notion that underwrites clarification interaction.

Thus, defining relevance involves interplay between semantic ontology, grammar, and interaction conventions. Various frameworks where relevance merely ties in the content of utterances (e.g. (Groenendijk and Roelofsen, 2009; van Benthem and Minica, 2009)) and even (Asher and Lascarides, 2003), which has admirably wide coverage, seem intrinsically unable to scale up to deal with metacommunicative relevance. Characterizing relevance requires a theory that allows ontology, grammar, and interaction to be encoded within the interaction conventions. For this purpose I employ Type Theory with Records (TTR) (Cooper, 2005). KoS and TTR are introduced in section 3. After which I sketch an attempt to combine the various notions of relevance so that they can be used to explicate examples of the type (1) above.

2 A Five Step Approach to Analyzing Relevance

2.1 Step 1: answerhood

In constructing our notion of relevance for queries, the first step is the most familiar. Any speaker of a given language can recognize, independently of domain knowledge and of the goals underlying an interaction, that certain propositions are *about* or *directly concern* a given question. This, I suggest, is the answerhood relation needed for characterizing interrogative relevance. It must be sufficiently inclusive to accommodate conditional, weakly modalized, and quantificational answers, all of which are pervasive in actual linguistic use, as in the following BNC examples:

(2) a. Christopher: Can I have some ice-cream then?
Dorothy: you can do if there is any. (BNC)

b. Anon: Are you voting for Tory?
Denise: I might. (BNC, slightly modified)

c. Dorothy: What did grandma have to catch?
Christopher: A bus. (BNC, slightly modified)

d. Elinor: Where are you going to hide it?
Tim: Somewhere you can't have it.

How to formally and empirically characterize aboutness is an interesting topic researched within work on the semantics of interrogatives (see e.g. (Ginzburg and Sag, 2000; Groenendijk, 2006)), though a comprehensive, empirically-based account is still elusive.

2.2 Step 2: q-specificity

The second step we take is somewhat less familiar and already a bit trickier. Any inspection of corpora, nonetheless, reveals the underdiscussed fact that many queries are responded to with a query. A

large proportion of these are clarification requests, to be discussed in section 2.5. But in addition to these, there are query responses whose content directly addresses the question posed, as exemplified in ((3)):

(3) a. A: Who murdered Smith? B: Who was in town?

b. A: Who is going to win the race? B: Who is going to participate?

c. Carol: Right, what do you want for your dinner?
Chris: What do you (pause) suggest? (BNC, KbJ)

d. Chris: Where's mummy?
Emma: Mm?
Chris: Mummy?
Emma: What do you want her for? (BNC, KbJ)

There has been much work on relations among questions within the framework of *Inferential Erotetic Logic* (IEL) (see e.g. (Wiśniewski, 2001; Wiśniewski, 2003)), yielding notions of q(uestion)–*implication*. From this a natural hypothesis can be made about such query responses, as in ((4))a); a related proposal, first articulated by (Carlson, 1983), is that they are constrained by the semantic relations of *dependence*, or its converse *influence*. A straightforward definition of these notions is in ((4)b). Its intuitive rationale is this: discussion of q_2 will necessarily bring about the provision of information about q_1:[4]

(4) a. q_2 can be used to respond to q_1 if $q_1 q-$ implies q_2. (Or q_2 influences q_1; Or q_1 depends on q_2)

b. q_2 influences q_1 iff any proposition p such that p Resolves q_2, also satisfies p entails r such that r is About q_1.

Question implication or dependence seem to constitute a sufficient condition for felicity of a (non-metacommunicative question) response. It does not seem to be a necessary condition. For instance, (3d) is felicitous but does not permit the inference

(5) Where Mummy is depends on what Chris wants for her.

[4]The definition of influence/dependence in ((4)b) makes reference to the answerhood notion of resolvedness, an agent–relative notion of exhaustiveness, as argued in (Ginzburg, 1995). Although for the moment I don't spell this out, this makes influence/dependence agent–relative rather than purely semantic notions, in contrast to aboutness. One could eliminate this asymmetry by using a purely semantic notion of exhaustiveness. This issue is further discussed below.

Consequently, a number of researchers have expressed doubt that it is dependence that underpins the requisite question/question (e.g. (Larsson, 2002; Shaheen, 2009a). Instead, with (e.g. (Asher and Lascarides, 2003)), they suggest that the requisite relation is plan–oriented, as could be articulated in terms of the rhetorical relation **Q(uery)-Elab(oration)** informally summarized in ((6)):

(6) If Q-Elab(α, β) holds between an utterance α uttered by A, where g is a goal associated by convention with utterances of the type α, and the question β uttered by B, then any answer to β must elaborate a plan to achieve g.

This latter proposal, motivated by interaction in cooperative settings, is vulnerable to examples such as ((7)):

(7) a. A: What do you like? B: What do you like?

b. A: What is Brown going to do about it? B: Well, what is Cameron?

I leave the precise characterization of this class of responses as an open issue, which requires more empirical research, both corpora-based and experimental, though for concreteness will assume an account based on q–implication/dependence.

2.3 Step 3: ability to answer

The first departure from a notion determined by questions *per se* is what one might call *metadiscursive* relevance. Irrelevance *implicatures* are an instance of metadiscursive interaction—interaction about what should or should not be discussed at a given point in a conversation:

(8) a. A: What's the problem with the drains?

b. B: I don't know.

c. B: You asked me that already.

d. B: You can't be serious.

e. B: Do we need to talk about this now?

f. B: I don't wish to discuss this now.

g. B: Whatever. Millie called yesterday.

I will mention one possible proposal concerning this aspect of relevance below. The crucial point metadiscursivity emphasizes is that a query introduces the potential for discussion of *other* questions. Specifically in this case the need to address the issue of whether a given question q should be discussed at a particular point by the responder B, an issue we might paraphrase informally as *?WishDiscuss(B, q)*.

2.4 Step 4: Genre specificity

In the case of metadiscursive relevance the issue introduced arises from the query interaction. Another source of relevance is the activity or genre type. Relevance driven by the domain plays an important role, as emphasized by a vast literature in AI, going back at least to (Cohen and Perrault, 1979; Allen and Perrault, 1980). In (9) B's perfectly relevant response is not *about* the query A asked:

(9) A: How can I help you?
B: A second class return ticket to Darlington, leaving this afternoon.

The basic intuition one can pursue is that a move can be made if it *relates to the current activity*. In some cases the activity is very clearly defined and tightly constrains what can be said. In other cases the activity is far less restrictive on what can be said:

(10) a. **Buying a train ticket**: c wants a train ticket: c needs to indicate where to, when leaving, if return, when returning, which class, s needs to indicate how much needs to be paid

b. **Buying in a boulangerie**: c needs to indicate what baked goods are desired, b needs to indicate how much needs to be paid

c. **Chatting among friends**: first: how are conversational participants and their near ones?

d. **Buying in a boulangerie from a long standing acquaintance**: combination of (b) and (d).

Trying to operationalize activity relevance presupposes that we can classify conversations into various *genres*, a term we use following (Bakhtin, 1986) to denote a particular type of interactional domain. There are at present remarkably few such taxonomies (though see (Allwood, 1999) for an informal one.) and we will not attempt to offer one here. However, as we will see below, we can indicate how to classify a conversation into a genre and build a notion of *genre-based* relevance from that.

2.5 Step 5: metacommunicative relevance

The final step for now will involve the most radical moves, ones that are ultimately difficult for many existing logical frameworks. In other words assessing which utterances are relevant as responses to an initial query—or any other type of move for that matter—requires reference to more than the query's content. This is demonstrated most clearly by metacommunicative responses, the two main types being acknowledgements of understanding and clarification requests (CRs). Here I mention a couple of salient facts that any account of metacommunicative relevance needs to address. First, CRs come in four main types, one of which relates to the phonological form of the utterance:

(11) A: Did Jo leave?

a. **intended content queries**:(*Jo?*),

b. **Repetition requests**: (*What?*),

c. **Relevance clarifications**: (*What do you mean?*),

d. **Requests for underlying motivation**: (*Why?*).

Second, there exist syntactic and phonological parallelism conditions on certain CR interpretations:

(12) a. A: Did Bo leave? B: Max? (cannot mean: intended content reading: **Who are you referring to?** or **Who do you mean?**)

b. A: Did he adore the book. B: adore? / #adored?

3 Relevance in KoS

As the underlying logical framework, I use Type Theory with Records (TTR) (Cooper, 2005), a model–theoretic descendant of Martin-Löf Type Theory (Ranta, 1994). What is crucial for current purposes about this formalism, which takes situation semantics as one of its inspirations, is that it provides access to both types and tokens at the object level. Concretely, this enables simultaneous reference to both utterances and utterance types, a key desideratum for modelling metacommunicative interaction. This distinguishes TTR from (standard) Discourse Representation Theory,[5] for instance, where the witnesses are at a model theoretic level, distinct from the level of discourse representations. The provision of entities at both levels of tokens and types allows one to combine aspects of the typed feature structures world and the set theoretic world, enabling its use as a computational grammatical formalism. The formalism can, consequently, be used to build a semantic ontology, and to write conversational interaction and grammar rules.

[5]There are versions of DRT that do allow for the presence of witnesses in the logical representation, e.g. Compositional DRT (Muskens, 1996), employed to underpin the PTT dialogue framework (Poesio and Rieser, 2010).

3.1 Information States

On the view developed in KoS, there is actually no single context, for reasons connected primarily with the integration of metacommunicative and illocutionary interaction, which I will touch on in section 3.5. Instead of a single context, analysis is formulated at a level of information states, one per conversational participant. The type of such information states is given in (13a). I leave the structure of the private part unanalyzed here, for details on this, see (Larsson, 2002). The dialogue gameboard represents information that arises from publicized interactions. Its structure is given in the type specified in (13b)——the *spkr,addr* fields allow one to track turn ownership, *Facts* represents conversationally shared assumptions, *Pending* and *Moves* represent respectively moves that are in the process of/have been grounded, *QUD* tracks the questions currently under discussion:

(13) a. TotalInformationState (TIS):

$$\begin{bmatrix} \text{dialoguegameboard} : \text{DGB} \\ \text{private} : \text{Private} \end{bmatrix}$$

b. DGB =

$$\begin{bmatrix} \text{spkr: Ind} \\ \text{addr: Ind} \\ \text{c-utt} : \text{addressing(spkr,addr)} \\ \text{Facts} : \text{set(Proposition)} \\ \text{Pending} : \text{list(locutionary Proposition)} \\ \text{Moves} : \text{list(locutionary Proposition)} \\ \text{QUD} : \text{poset(Question)} \end{bmatrix}$$

Context change is specified in terms of *conversational rules*, rules that specify the *effects* applicable to a DGB that satisfies certain *preconditions*. This allows both illocutionary effects to be modelled (preconditions for and effects of greeting, querying, assertion, parting etc), interleaved with *locutionary effects*. How querying works in this framework I will illustrate in the next section, once we have discussed *q(uestion)–specificity*.

3.2 Questions in context

The basic notion of relevance that has emerged so far can be summarized in term of the notion of q-specificity in ((14)):

(14) q-specific utterance: an utterance whose content is either a proposition p About q or a question q_1 on which q Depends

This can be embedded in 2-person interaction via a protocol as in ((15)):

(15)

querying	assertion
LatestMove = Ask(A,q)	LatestMove = Assert(A,p)
A: push q onto QUD; release turn;	A: push p? onto QUD; release turn
B: push q onto QUD; take turn; make q—specific utterance take turn.	B: push p? onto QUD; take turn; Option 1: Discuss p? Option 2: Accept p
	LatestMove = Accept(B,p)
	B: increment FACTS with p; pop p? from QUD;
	A: increment FACTS with p; pop p? from QUD;

As argued in (Ginzburg, 2011), the only query specific aspect of the query protocol in (15) is the need to increment QUD with q as a consequence of q being posed:

(16) Ask QUD–incrementation:

$$\begin{bmatrix} \text{pre} : \begin{bmatrix} \text{q} : \text{Question} \\ \text{LatestMove} = \text{Ask(spkr,addr,q):IllocProp} \end{bmatrix} \\ \text{effects} : \left[\text{qud} = [\text{q,pre.qud}] : \text{list(Question)}\right] \end{bmatrix}$$

The specification `make q-specific utterance` is an instance of a general constraint that characterizes the contextual background of reactive queries and assertions. This specification can be formulated as in ((17)): the rule states that if q is QUD–maximal, then either participant may make a q–specific move. Whereas the preconditions simply state that q is QUD–maximal, the preconditions underspecify who has the turn and require that the latest move—the first element on the MOVES list—stand in the *Qspecific* relation to q:[6]

(17) QSpec

$$\begin{bmatrix} \text{preconds} : \left[\text{qud} = \langle \text{q, Q} \rangle : \text{poset(Question)}\right] \\ \text{effects} : \begin{bmatrix} \text{spkr} : \text{Ind} \\ \text{c1} : \text{spkr} = \text{preconds.spkr} \vee \text{preconds.addr} \\ \text{addr} : \text{Ind} \\ \text{c2: member(addr,}\{\text{preconds.spkr,preconds.addr}\}\text{)} \\ \wedge\ \text{addr} \neq \text{spkr} \\ \text{r} : \text{AbSemObj} \\ \text{R: IllocRel} \\ \text{Moves} = \langle \text{R(spkr,addr,r)} \rangle \bigoplus m : \text{list(IllocProp)} \\ \text{c1} : \text{Qspecific(r,preconds.qud.q)} \end{bmatrix} \end{bmatrix}$$

The notion of q–specificity still needs some refinements if it is to do its job of regulating responses that address a given question. The most direct refinement concerns *indirect answerhood*:

[6] This underspecification of turn ownership is the basis for a unified account of question posing in monologue, 2-person querying, and multilogue provided in (Ginzburg, 2011).

responses that provide an answer indirectly should clearly be accommodated (Asher and Lascarides, 1998). This means relativizing aboutness by an entailment notion based on common ground information represented in FACTS.

3.3 Metadiscursive Relevance

A natural way to analyze such utterances is along the lines of the conversational rule QSPEC discussed in section 3.2: A introducing q gives B the right to follow up with an utterance about an issue we could paraphrase informally as *?WishDiscuss(q)*. Such a CCUR is sketched in ((18)):

(18) Discussing u?

$$\begin{bmatrix} \text{preconds} & : \text{DGB} \\ \text{effects} & : \begin{bmatrix} \text{spkr = preconds.addr : Ind} \\ \text{addr = preconds.spkr : Ind} \\ \text{r : AbSemObj} \\ \text{R: IllocRel} \\ \text{Moves} = \langle \text{R(spkr,addr,r)} \rangle \\ \bigoplus \text{pre.Moves : list(IllocProp)} \\ \text{c1 : Qspecific(R(spkr,addr,r),} \\ \text{?WishDiscuss(pre.maxqud)} \\ \text{qud} = \left\langle \begin{matrix} \text{?WishDiscuss(pre.maxqud),} \\ \text{pre.qud} \end{matrix} \right\rangle \\ \text{: poset(Question)} \end{bmatrix} \end{bmatrix}$$

3.4 Genre-based Relevance

An account of genre-based relevance presupposes a means of classifying a conversation into a genre.[7] One way of so doing is by providing the description of an information state of a conversational participant who has *successfully* completed such a conversation. Final states of a conversation will then be records of type T for T a subtype of DGB_{fin}, here Questions No (longer) Under Discussion (QNUD) denotes a list of issues characteristic of the genre which will have been resolved in interaction:

(19) $\text{DGB}_{fin} = \begin{bmatrix} \text{Facts : Prop} \\ \text{QNUD = list : list(question)} \\ \text{Moves : list(IllocProp)} \end{bmatrix}$

In ((20)) we exemplify two genres, informally specified in (10):

(20) a. CasualChat:

$$\begin{bmatrix} \text{A, B : Ind} \\ \text{t: TimeInterval} \\ \text{c1 : Speak(A,t)} \vee \text{Speak(B,t)} \\ \text{facts : Set(Prop)} \\ \text{qnud : list(question)} \\ \text{c2: } \{\lambda P.P(A), \lambda P.P(B)\} \subset \text{qnud} \\ \text{moves : list(IllocProp)} \end{bmatrix}$$

b. BakeryChat:

$$\begin{bmatrix} \text{A, B : Ind} \\ \text{t: TimeInterval} \\ \text{c1 : Speak(A,t)} \vee \text{Speak(B,t)} \\ \text{facts : Set(Prop)} \\ \text{qnud : list(question)} \\ \text{c2: } \left\{ \begin{matrix} \lambda P.P(A), \lambda P.P(B), \\ \lambda x.\text{InShopBuy(A,x)}, \\ \lambda x.\text{Pay(A,x)} \end{matrix} \right\} \subset \text{qnud} \\ \text{moves : list(IllocProp)} \end{bmatrix}$$

We can then offer the following definition of *activity relevance*: one can make a move m0 if one believes that that the current conversation updated with m0 is of a certain genre G0. Making move $m0$ given what has happened so far (represented in $dgb0$) can be *anticipated* to conclude as final state $dgb1$ which is a conversation of type G0:

(21) m0 is relevant to G0 in dgb0 for A iff A believes that there exists dgb1 such that $(dgb0 \bigoplus m0) \sqsubset dgb1$, and such that dgb1 : G0

3.5 Metacommunicative Relevance

In the immediate aftermath of a speech event u, **Pending** gets updated with a record of the form $\begin{bmatrix} \text{sit = u} \\ \text{sit-type = T}_u \end{bmatrix}$ (of type *locutionary proposition* (LocProp)). Here T_u is a grammatical type for classifying u that emerges during the process of parsing u. The relationship between u and T_u—describable in terms of the proposition $p_u = \begin{bmatrix} \text{sit = u} \\ \text{sit-type = T}_u \end{bmatrix}$— can be utilized in providing an analysis of grounding/CRification conditions:

(22) a. Grounding: p_u is true: the utterance type fully classifies the utterance token.

b. CRification: p_u is false, either because T_u is weak (e.g. incomplete word recognition) or because u is incompletely specified (e.g. incomplete contextual resolution).

[7] For an application of genre-based relevance to the semantics of *Why*-questions, see (Shaheen, 2009b).

In case p_u is true, p_u becomes the LatestMove and relevance possibilities discussed above come into operation. Otherwise clarification interaction ensues. This involves accommodation of questions into context by means of a particular class of conversational rules—Clarification Context Update Rules (CCURs), whose general substance is paraphrased in ((23)a), with a particular instance given in ((23)b):

(23) a. CCUR$_i$: given u1 a constituent of MaxPending, accommodate as MaxQUD q$_i(u1)$, follow this up with an utterance which is *co-propositional* with q$_i(u1)$.

b. Parameter identification: Input:
$$\begin{bmatrix} \text{Spkr : Ind} \\ \text{MaxPending : LocProp} \\ \text{u1} \in \text{MaxPending.sit.constits} \end{bmatrix}$$
Output:
$$\begin{bmatrix} \text{MaxQUD = What did spkr mean by u1?} \\ \text{LatestMove : LocProp} \\ \text{c1: CoProp(LatestMove.cont,MaxQUD)} \end{bmatrix}$$

CoPropositionality for two questions means that, modulo their domain, the questions involve similar answers. For instance 'Whether Bo left', 'Who left', and 'Which student left' (assuming Bo is a student.) are all co-propositional:

(24) a. Two utterances u_0 and u_1 are *co-propositional* iff the questions q_0 and q_1 they contribute to QUD are co-propositional.

(i) qud-contrib(m0.cont) is m0.cont if m0.cont : Question

(ii) qud-contrib(m0.cont) is ?m0.cont if m0.cont : Prop[8]

b. q_0 and q_1 are co-propositional if there exists a record r such that $q_0(r) = q_1(r)$.

In the current context co-propositionality amounts to: either a CR which differs from MaxQud at most in terms of its domain, or a correction—a proposition that instantiates MaxQud.

4 Combining Relevance

What then does Relevance amount to? Pretheoretically, Relevance relates an utterance u to an information state I just in case there is a way to successfully update I with u. Let us restrict attention for now to the case where the input context is a query. Given a set of conversational rules $\mathcal{C}$, a grammar $\mathcal{G}$ and an information state $I_0 : TIS$, an utterance u is $\mathbf{U}(\mathbf{tterance})_{\mathcal{C},\mathcal{G}}{}^{I_0}$**-relevant** iff there exist $c_1, \ldots, c_{k+1} \in \mathcal{C}, T_u \in \mathcal{G}, k \geq 0$ such that $c_1(I_0) = I_1, \ldots, c_{k+1}(I_k) = I_{k+1}$, where C's information state I_0 satisfies ((25)a); where by means of a sequence of updates the locutionary proposition $p_u = prop(u, T_u)$ becomes the value of LatestMove (condition ((25)b); and the final element of the sequence of updates I_{k+1} is such that one of the conditions in ((25)c-f) is satisfied—u is either q–specific, an appropriate CR, relates to the issue of willingness to discuss q, or is genre–relevant:

(25) a. I_0.DGB.LatestMove = v; v.content = Ask(A,q),

b. $I_{k+1}.DGB.LatestMove = p_u$

c. $p_u.content$ is q–specific relative to I.DGB, Or

d. $p_u.content$ is CoPropositional with some question q0 that satisfies q0 = CCUR1.effects. maxqud($I_0.DGB.MaxPending$) for some Clarification Context Update Rule CCUR1, Or

e. $p_u.content$ is q_0–specific, where q_0 is the question ?WishDiscuss(B,q), Or

f. One of C's beliefs in I_0 is that: for some G0 there exists dgb1 such that $(I_0.DGB \bigoplus p_u) \sqsubset dgb1$, and such that dgb1 : G0

A number of remarks can be made about (25), primarily about the relata of this notion.

- The definition is relative to both the set of conversational rules and to a grammar from which the types T_u from which locutionary propositions originate.
- Relevance is, by and large, DGB oriented. Only ((25)f) explicitly involves reference to the entire information state.

5 Using Relevance

In this section I offer one application of the internalized notion of relevance, formulating a rule underwriting lack of wish to address an utterance.[9] A prototypical example in this respect is given in ((26)a). Two further examples from literary texts convey a similar import:

(26) A: Horrible talk by Rozzo. B: It's very hot in here. **Implicates**: B does not wish to discuss A's utterance.

[8] Recall from the assertion protocol that asserting p introduces p? into QUD.

[9] In seeking to underwrite this inference via conversational rule there is no inconsistency with a Gricean view that such an implicature can be explicated in terms of calculations made by rational agents on the basis of apparent violations of the Cooperative Principle etc. This rule represents a "short circuited" version of the Gricean account.

a. Rumpole: Do you think Prof Clayton killed your husband? Mercy Charles: Do you think you'll get him off? ('Rumpole and the Right to Silence', p. 100)

b. Harry: Is that you James? Stella: What? No, it isn't. Who is it? Harry: Where's James? Stella: He's out. Harry: Out? Oh, well, all right. I'll be straight round. Stella: What are you talking about? Who are you? (Pinter, *The Collection*, p. 133)

In current terms we could formulate the inference as in ((27)):

(27) $\neg$Relevant(u,I) $\mapsto$
A does not wish to address I.dgb.LatestMove.

More formally, we can offer the update rule in ((28))—given that MaxPending is *irrelevant* to the DGB, one can make MaxPending into LatestMove while updating Facts with the fact that the speaker of MaxPending does not wish to discuss MAX-QUD:

(28)
$$\left[\begin{array}{l}\text{preconds: } \left[\begin{array}{l}\text{I : TIS}\\ \text{c: } \neg\text{Relevant(maxpending,I)}\end{array}\right]\\ \text{effects : } \left[\begin{array}{l}\text{LatestMove = pre.pending : LocProp}\\ \text{Facts = pre.Facts} \cup\\ \left\{\neg\ \text{WishDiscuss(pre.spkr,pre.maxqud)}\right\}.\end{array}\right]\end{array}\right]$$

Note that this does not make the *unwillingness to discuss* be the *content* of the offending utterance; it is merely an inference. Still this inference will allow MAX-QUD to be downdated from the DGB via the general mechanisms that regulate QUD downdate in conjunction with FACTS update.

6 Conclusions

Relevance is the most fundamental notion for research on dialogue. Restricting attention here to the case where a query has just taken place, I have elucidated four dimensions of this notion: responses that are question–specific, metadiscursive, genre–specific, and metacommunicative. I have formalized this notion, starting with the intuition that it amounts to a relation between an utterance and an information state where the utterance can successfully update the information state whose most recent move is a query. using the dialogue theory KoS and the formalism of Type Theory with Records. The notion of relevance that emerges is primarily one grounded in publicized contextual information, though it has some important unpublicized components, primarily those relating to genre–dependent knowledge. I have motivated the need for a notion of relevance internalized in the theory of meaning via its application in a class of clarification requests ('What do you mean') and the celebrated Gricean implicatures of lack of desire to address an utterance.

KoS enables us to construct a potentially rich theory of relevance. But as I have made clear there is a slew of issues we are in the dark about. These include:

1. **Empirical coverage**: what aspects does the four cornered characterization offered above intrinsically miss?

2. **The nature of q-responsiveness**: is there a clean way, analogous to answerhood, to characterize the (non-metacommunicative) questions arising from a given query?

3. **Relevancial success**: (Why) are there in practice few relevance CRs?

Acknowledgments

I would like to thank Jon Shaheen for a series of illuminating comments on an extended version of this paper and Robin Cooper and Alex Lascarides for discussion of topics discussed therein.

References

James Allen and Ray Perrault. 1980. Analyzing intention in utterances. *Artificial Intelligence*, 15:143–178.

Jens Allwood. 1999. The swedish spoken language corpus at göteborg university. In *Proceedings of Fonetik 99*, volume 81 of *Gothenburg Papers in Theoretical Linguistics*.

Nicholas Asher and Alex Lascarides. 1998. Questions in dialogue. *Linguistics and Philosophy*, 21.

Nicholas Asher and Alex Lascarides. 2003. *Logics of Conversation*. Cambridge University Press, Cambridge.

M.M. Bakhtin. 1986. *Speech Genres and Other Late Essays*. University of Texas Press.

Lauri Carlson. 1983. *Dialogue Games*. Synthese Language Library. D. Reidel, Dordrecht.

Philip Cohen and Ray Perrault. 1979. Elements of a plan-based theory of speech acts. *Cognitive Science*, 3:177–212.

Robin Cooper. 2005. Austinian truth in martin-löf type theory. *Research on Language and Computation*, pages 333–362.

Raquel Fernández. 2006. *Non-Sentential Utterances in Dialogue: Classification, Resolution and Use*. Ph.D. thesis, King's College, London.

Jonathan Ginzburg and Robin Cooper. 2004. Clarification, ellipsis, and the nature of contextual updates. *Linguistics and Philosophy*, 27(3):297–366.

Jonathan Ginzburg and Raquel Fernández. 2010. Dialogue. In Alex Clark, Chris Fox, and Shalom Lappin, editors, *Handbook of Computational Linguistics and Natural Language*, Oxford. Blackwell.

Jonathan Ginzburg and Ivan A. Sag. 2000. *Interrogative Investigations: the form, meaning and use of English Interrogatives*. Number 123 in CSLI Lecture Notes. CSLI Publications, Stanford: California.

Jonathan Ginzburg. 1994. An update semantics for dialogue. In H. Bunt, editor, *Proceedings of the 1st International Workshop on Computational Semantics*. ITK, Tilburg University, Tilburg.

Jonathan Ginzburg. 1995. Resolving questions, i. *Linguistics and Philosophy*, 18:459–527.

Jonathan Ginzburg. 2011. *The Interactive Stance: Meaning for Conversation*. Oxford University Press, Oxford. Draft available from http://www.dcs.kcl.ac.uk/staff/ginzburg/tis1.pdf.

J. Groenendijk and F. Roelofsen. 2009. Inquisitive semantics and pragmatics. In *Meaning, Content, and Argument: Proceedings of the ILCLI International Workshop on Semantics, Pragmatics, and Rhetoric. www. illc. uva. nl/inquisitive-semantics*.

Jeroen Groenendijk. 2006. The logic of interrogation. In Maria Aloni, Alistair Butler, and Paul Dekker, editors, *Questions in Dynamic Semantics*, volume 17 of *Current Research in the Semantics/Pragmatics Interface*, pages 43–62. Elsevier, Amsterdam. An earlier version appeared in 1999 in the Proceedings of SALT 9 under the title 'The Logic of Interrogation. Classical version'.

Staffan Larsson. 2002. *Issue based Dialogue Management*. Ph.D. thesis, Gothenburg University.

R. Muskens. 1996. Combining Montague semantics and discourse representation. *Linguistics and Philosophy*, 19(2):143–186.

Massimo Poesio and Hannes Rieser. 2010. (prolegomena to a theory of) completions, continuations, and coordination in dialogue. *Dialogue and Discourse*, 1.

M. Purver. 2006. Clarie: Handling clarification requests in a dialogue system. *Research on Language & Computation*, 4(2):259–288.

Aarne Ranta. 1994. *Type Theoretical Grammar*. Oxford University Press, Oxford.

Jon Shaheen. 2009a. Comments on Ginzburg's paper, i: Intuitions about query responses and hesitation marking. Presented at the Michigan Fall Symposium on Linguistics and Philosophy, University of Michigan.

Jon Shaheen. 2009b. Genre-based relevance and why-questions. University of Amsterdam unpublished Ms.

A.M. Turing. 1950. Computing Machinery and Intelligence. *Mind*, 59(236):433–460.

Johan van Benthem and Stefan Minica. 2009. Toward a dynamic logic of questions. In Xiangdong He, John Horty, and Eric Pacuit, editors, *Proceedings of Logic, Rationality and Interaction (LORI-II)*.

Andrzej Wiśniewski. 2001. Questions and inferences. *Logique et Analyse*, 173:5–43.

Andrzej Wiśniewski. 2003. Erotetic search scenarios. *Synthese*, 134:389–427.