

Pragmatics and Translation

David Farwell

Stephen Helmreich

Computing Research Laboratory

New Mexico State University

Las Cruces, New Mexico

In this paper, we present an approach to machine translation rooted in pragmatic¹. It is an approach which relies on:

- the use of context to interpret source language utterances and to produce target language correlates,
- a context of the utterance which crucially includes nested beliefs environments which are constructed and modified through ascription during processing,
- knowledge of the world which is accessed for constructing or modifying the context,
- context-sensitive (non-monotonic) inferencing within the context to resolve ambiguities during interpretation or to select expressions during production.

In Section 1, we will begin the discussion by motivating the need for a pragmatics-based approach which takes into account the beliefs of the participants in the translation process by focussing on one example. In Section 2, we show that such examples are not isolated instances but pervade translations. We show that there are patterns of differences in multiple translations of the same text which are related to different global interpretations of what the text itself is about or to differing world views on the part of the translators. In Section 3, we present a number of concepts related to translation from a pragmatics-based perspective and suggest a

1. Work on this project was supported under DOD grant: MDA904-92-C-5189.

possible computational framework for implementing such an approach. In the final section, we conclude by discussing the implications of a pragmatics-based approach for translation analysis, translation evaluation and for future directions in machine translation research.

This discussion represents a summary of work presented in more detail in Farwell and Helmreich (1993, 1995, 1996, 1997 and 1998).

Section 1

In this section, we argue for a pragmatics-based approach to machine translation. The need for pragmatics in Natural Language Processing and Machine Translation has long been recognized (Bar-Hillel 1960, Wilks 1975, Nirenburg *et al.* 1992). Yet pragmatics (as the study of language in context) is often difficult to distinguish from semantics (the study of the connection between the language sign system and the world it represents)². The history of Machine Translation shows a progression towards systems containing more and more knowledge in order to represent the meaning of the source text. We take as a

2. Many theorists such as Pustejovsky (1991), Lakoff (Lakoff & Johnson 1980, Lakoff 1987), or Nirenburg (1987) appear to us to coalesce language meaning with language use and thus are interested in semantics. Other theorists such as Fauconnier (1985), Green (1984), or Nunberg (1978), place emphasis on the issues of language use and appear to us to attempt to subsume semantics within pragmatics.

point of departure the Pangloss/Mikrokosmos project which represents the most consistent, large-scale attempt to develop a knowledge-based, interlingual machine translation system to date (Pangloss 1994).

As might be expected, various pragmatic issues have been addressed within the Pangloss system. The Mikrokosmos MT engine, for example, monitors and uses information about speakers and hearers, discourse structure, and world knowledge in constructing its interlingual representation of an input text (Attardo 1994, Carlson & Nirenburg 1990). It can handle metonymic constructions and metaphors (Onyshkevych & Nirenburg 1994). However, we suggest that it is still a semantics-based system. Even though contextual information and world-knowledge are both represented and used inferentially, the goal is to produce a semantic representation of the text.

Language, however, is used not simply to report events in the world. It is also used to convey the rich mental models that individuals and cultures bring to bear on the communication process. It is the claim of a pragmatics-based approach that texts do not have meanings, but rather that in producing texts, people intend meanings. Thus, a text can only be approached through an interpretation. That is to say, the translator attempts to understand the author's intent in creating the source text for the original audience and then recreates, to the extent possible, that intent for the target audience using the target language³.

3. We do not suggest this as a psychological model of translation. Translation, we believe, is an art or skill and not a faculty (in the sense of Fodor 1984) and can therefore be approached in any number of different ways. We suggest only that this is a plausible and appropriate model for computational representation.

As is obvious here and throughout this article, the influence of the work of Grice (1975) is clear. Implicatures and inferences from "what is said" are as fully part of the communication as the semantic content of the utterance itself.

This pragmatics-based approach, then, focuses on providing interpretations of a text that insure a coherent account of the intent behind the text. Such an approach relies heavily on representations of the beliefs and other mental attitudes (such as expectations, hopes, likes, and dislikes) of the participants in the communicative process: the author, the translator, and their respective audiences (the addressees).⁴ In particular, then, we require a system that can model:

- the translator's beliefs,
- the translator's beliefs about the author's beliefs,
- the translator's beliefs about the author's beliefs about the (source language) addressee's beliefs,
- the translator's beliefs about the (target language) addressee's beliefs.

Such system models can be found in the work on ViewGen (Ballim & Wilks 1991), a system which constructs beliefs spaces for any number of agents about any number of topics. In addition, we find a computational basis for the required inferencing in the work associated the ATT-Meta system (Barnden *et al.* 1994). This system includes a default inference mechanism, an epistemic logic involving four possible positive truth conditions (certain, default, probable, and possible), a truth maintenance system, metaphorical pretense cocoons (for treating metaphors as true), and simulative reasoning (for inferencing within embedded belief states).

Motivating Pragmatics-based Translation

For the sake of exposition within a limited space, we motivate the need for a pragmatics-based approach by way of a case involving the translation of a single

4. In this paper we will use the term "belief" to refer to any epistemic state. We do not distinguish (except by degree of commitment) between beliefs, knowledge, hypotheses, and thoughts. For us, they are all beliefs, though of course, held with different degrees of tenacity.

French noun phrase, taken from a subtitled motion picture, *Jesus of Montreal* (Arcand 1989). We have devised all the alternative translations except for the original subtitle. We discuss three alternative translations which stem from (a) alternative views of the beliefs of the film's audience (the addressees of the translation) and (b) alternative views of the translator's beliefs about the world.

We will simplify matters by assuming that the translator is sitting beside us interpreting the film as it develops, each character an independent agent. Thus, for any given utterance, there are four relevant participants: the actor who speaks, the actor who is addressed, the translator and the audience (the addressees of the translation). We also ignore the additional complexities involved in the fact that both actors are also speaking in some sense on behalf of the author of the screen play to the film's original audience.

As background (i.e., that part of the context of which the protagonists, the translator and the film's audience are all aware), we provide the following synopsis.

A priest at a shrine outside Montreal has been sponsoring a religious drama every summer for 35 years. Since the text has become somewhat outdated, he asks Coulombe, a young actor who has recently returned from an extended sojourn, to modernize the script and to play the part of Jesus. He agrees and immediately sets about looking for collaborators. The priest suggests that Constance, an old friend of Coulombe's, would be a good person for Coulombe to enlist in his endeavor and so he seeks her out. She agrees to work with him, and, in passing, invites him to stay at her apartment. He agrees.

In a later scene, the scene of our attention, Coulombe returns earlier than usual to what he assumes is an empty apartment. He starts to make himself comfortable, making some noise in the process. At this point he hears someone moving about in Constance's

bedroom and, suddenly, she emerges from within, closing the door behind her. She says, *T'es déjà là, toi?* (Back already?) and then, coughing significantly, says to herself, *Bon...* (Okay...). At this point Coulombe realizes that there may be someone else in the bedroom and whispers, *Tu veux que je m'en aille?* (Should I go?). She shakes her head no, laughs nervously, opens the door and says to whomever is inside, *Ben, écoutes, sors* (Come on out), *On va pas jouer une scène de Feydeau* (This isn't a bedroom farce).

It is this last utterance and its subtitle that we wish focus on. The translator who provided the subtitles for the film has glossed *On va pas jouer une scène de Feydeau* as *This isn't a bedroom farce*. This is not the only possible translation, of course, and we look at alternatives so as to draw out the underlying assumptions that determined the translator's choices.

At the time of Constance's utterance, the protagonists, the translator and the audience have the following beliefs (among others).

Coulombe is living in Constance's apartment.

They are clearly close friends and colleagues.

Coulombe has entered the apartment unexpectedly early.

It is still mid-afternoon.

He accidentally makes a loud noise.

Constance emerges from her bedroom dressed in a nightgown and closes the door behind her.

She is somewhat flustered by Coulombe's unexpected presence.

Coulombe believes there is someone else in Constance's bedroom that he has caught them in a compromising situation.

Coulombe believes that Constance and the other person might prefer some privacy.

Coulombe believes that if Constance and the other person might wish to keep the identity of the other secret.

Constance believes that Coulombe believes that she and the other have

been caught in a compromising situation.

Constance believes that Coulombe believes that they might prefer some privacy.

Constance wishes to change Coulombe's belief.

Constance tells Coulombe not to leave.

Constance tells the person in the bedroom to come out and show him/herself.

Without going into the details, the analysis of the utterance begins by establishing, on linguistic grounds, that Constance is using *on va pas jouer* to express to the person in the bedroom that she does not wish the current situation (such as Coulombe's discovery of her and the unknown person alone together in her bedroom) to be understood as being of a to-be-determined type of play (i.e., *we're not playing ...*, *we're not going to play ...*). The next step is to assign an interpretation to *un scène de Feydeau*. Again, on linguistic grounds, coupled in this case with knowledge of the world, we establish that Constance is using *un scène de Feydeau* to refer to the type of situation that might be used as a scene in a play by the 19th century French playwright Feydeau who wrote bedroom farces (i.e., *a scene from a bedroom farce*). To arrive at this interpretation, it must be the case that:

Constance believes Feydeau is a playwright and that Feydeau wrote bedroom farces.

Constance believes the person in her bedroom believes Feydeau is a playwright and that Feydeau wrote bedroom farces.

The interpretation is completed by confirming that the situation under discussion (i.e., Constance and someone alone together in her bedroom) is indeed one that Feydeau may have written about. This becomes especially plausible when it turns out that the man in Constance's room is a priest, in fact, the very same priest who hired Coulombe to update the play.

Having arrived at an interpretation, the translator now needs to provide an equivalent expression for an English speaking audience. To express that some current situation is not of some type, he/she selects the expression *This is no ...* or *This isn't ...* or some such English equivalent. As for a situation typical of a bedroom farce of the sort that Feydeau might write about, e.g., two people getting caught in a compromising position by a significant other,⁵ the translator checks his/her beliefs about the audience. If it is assumed that the addressee of the translation would not typically believe that Feydeau is a playwright or that Feydeau wrote bedroom farces, quite possible for those unfamiliar with French culture or with the theater, then reference to Feydeau will fail to have the intended effect and some alternative expression must be chosen, e.g., *a bedroom farce* or *a scene from a bedroom farce*.

This leads to the first case of variation in translation stemming from variations in the translator's beliefs, namely, those based on variations in the beliefs the translator attributes to the addressee of the translation, the non-French speaking audience of the film. In the event that the translator assumes that the film's audience has the same beliefs about Feydeau as the speaker (Constance) and the addressee (the unknown person in the bedroom), he/she would most likely take advantage of those beliefs to provide a translation that more closely approximates the source language utterance in form and content, relying on the addressees of the

5. There appear to be two scenarios that could be drawn from a bedroom farce. One involves, as suggested above, being caught in a compromising situation by a jealous husband or lover and thus enduring an unpleasant scene. The other involves being apprehended by a moral or social arbiter, so that some unacceptable behavior is exposed. These two results: jealous rage or exposure, usually result (at least temporarily in bedroom farces) in public shame and humiliation, and in possible long-term negative social consequences such as divorce or loss of position.

translation to use those beliefs appropriately to interpret Constance's utterance. That is, if the translator assumes the addressees of the translation believe that Feydeau is a French playwright who wrote bedroom farces, then he/she would most likely have glossed the utterance, on the basis of the rhetic act,⁶ as *This is not a scene from Feydeau*. In fact, it appears that the translator assumes that the non-French speaking audience of the film has beliefs about Feydeau different from the beliefs of Constance and the unknown person in the bedroom, or, perhaps, has no beliefs about Feydeau at all. Thus, the translator avoids any expression that would rely on such beliefs for the interpretation of the translation and glosses the utterance, on the basis of the illocutionary act, as *This isn't a bedroom farce*.

6. We distinguish three levels of interpretation, which we identify loosely with three of Austin's speech acts (Austin 1962). At the first level, the intent of the speaker is to utter words of a language with a more or less definite sense or reference (p. 95). This is the rhetic act, and it is interpreted correctly if the hearer can identify the senses and references intended by the speaker. This is essentially what is normally called the semantic content of the utterance. At a second level of interpretation, the hearer attempts to understand the actual message that the speaker wishes to communicate, which may or, more likely, is not identical with the semantic content, but rather the result of the interference from this content and other premises drawn from assumed knowledge or context. We identify this loosely with the illocutionary act, the communicative intent of the utterance. Finally, at the broadest level, there is an intent to produce a change in the hearer. This change is the perlocutionary effect of the utterance, and is an interpretation of what the hearer understands to be the purpose or goal of the speaker in making the utterance. In the utterance at hand, the semantic content (rhetic act) is that *we are not playing a scene from Feydeau*. The communicative content (illocutionary act) is to tell the addressee that the current situation is not one in which he need fear disclosure of his identity. The perlocutionary intent of the utterance is to encourage the addressee to come out of the bedroom by assuring him that it is safe to do so.

A second case of translation variation based on variations in the beliefs-context concerns variations in the beliefs of translators about the world. It is possible, for instance, that the translator did not have the necessary beliefs about Feydeau to work out the intended interpretation. This is, of course, the bane of translators: lack of the relevant knowledge of the world, and it is far more common than appreciated even in extremely mundane discourse such as general news articles. It need not, however, deter translation, even appropriate translation. When there is a lack of knowledge, the translator must fall back to an even more abstract level of interpretation as a basis for the translation, that of the functional effect of the utterance. (Or, of course, the translator could fall back to a translation of the rhetic act, hoping that his addressees will be able to interpret what he cannot.) In our case, Constance is informing the unknown person in the bedroom that Coulombe will not be scandalized by their liaison and that the unknown person in the bedroom can safely show himself. This may be accomplished by glossing *On va pas jouer une scène de Feydeau*, on the basis of the intended perlocutionary effect of the utterance as *There's nothing to worry about; it's safe*. This strategy, however, can be rather dangerous since there is little data to aid in identifying the specific functional effect.

Summary

We believe that this example demonstrates that an approach to language rooted in pragmatics offers significant advantages over an approach that sees pragmatics as simply an additional component of the language system. In applying this approach to the field of language translation, we have shown how such an approach can provide explanations of the many possible translations for the same text.

Section 2

In this section we present certain concepts related to a pragmatics-based approach to Machine Translation and sketch out an informal processing model. We begin by developing notions of discourse context and utterance context and then outline a two-stage translation process involving interpretation and translation. Next, we introduce the notion of "user-friendly translation" and, finally, conclude with a discussion of translation equivalence.

Background

As a context for the discussion, we consider two differing translations into English of *el tercer piso* and *el segundo piso* in the following Spanish sentence taken from a news article about the Moscow real estate market in the early 1990's.

... los 300 metros cuadrados del **tercer piso** estaban disponibles pero fueron aguilados ..., sólo queda **el segundo piso**

While one translator has rendered these expressions as *the third floor* and *the second floor* respectively, another has rendered them as *the fourth floor* and *the third floor*. Although these two translations are clearly different, they are, in fact, both accurate and they are not necessarily logically inconsistent. The reason resides in the differing beliefs the translators have about the beliefs of the author and addressees of the Spanish text, and of the audience of the English translation.

We hypothesize that first translator assumes that the author of the text shares the translator's floor naming convention (say, using *ground floor*, *first floor*, etc. for the levels of a building as opposed to *first floor*, *second floor*, etc.) and that the addressees of the translation also share the translator's floor naming convention (though it is also possible that the author and addressees share a convention which the translator does not). Thus, the first translator refers to

the fourth level above ground as *the third floor* and the third level above ground as *the second floor*. If those assumptions are correct (and we do not have the crucial information to determine this), then the first translator's translation is equivalent at both the level of the rhetic (semantic) and illocutionary (communicative) actst. Otherwise, the translation will be less equivalent since the addressees of the translation will have to access information about alternative floor naming conventions and make the appropriate inferences in order to arrive at the author's intended meaning.

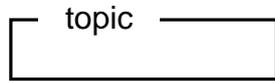
We hypothesize that the second translator assumes that either the author does not share the translator's floor naming convention (say, using *first floor*, *second floor*, etc. vs *ground floor*, *first floor*, etc.) or, alternatively, the addressees of the translation do not share the author's floor naming convention. Thus, the second translator refers to the fourth level above ground as *the fourth floor* and the third level above ground as *the third floor*. If either of those sets of assumptions is correct (again, we do not have the crucial information to determine whether they are), then the second translator's translation is equivalent at least at the illocutionary level of communicative content. Otherwise, the translation will be less equivalent since the addressees of the translation will have to access information about alternative floor naming conventions and make the appropriate inferences in order to arrive at the author's intended meaning.

Beliefs and Inferencing

As a framework for modeling the translation process, we adopt a beliefs ascription mechanism (such as that of Ballim & Wilks 1991) for constructing the relevant, recursively embedded beliefs spaces of the participants in the translation, ontologies (such as those of Nirenburg *et al.* 1995) for representing the beliefs within these spaces, and a default inferencing

engine (such as that of Barnden *et al.* 1994) for carrying out the inferencing over these beliefs within these spaces.

Following Ballim & Wilks 1991, beliefs about a topic are represented as a box labeled in the upper left-hand corner.



Beliefs of an agent about a topic are represented as a box labeled in the lower center which, in turn, contains a topic box.



The basic default rule for ascribing beliefs is:

unless there is specific evidence to the contrary, agent₁ ascribes its beliefs to agent₂.

Evidence to the contrary consists of preexisting beliefs of the target agent that are contradictory to or inconsistent with the beliefs being ascribed. We indicate the application of this rule by an arrow pointing from the source agent's environment to the target agent's environment. In Figure 1,

Agent₁'s beliefs about a topic have been ascribed to Agent₂.

Following Barnden *et al.* 1994, we further assume the existence of default (defeasible) inferencing engine of the sort used in ATT-Meta. This reasoning mechanism, when given a goal to prove, will evaluate all evidence chains that it can find for both the goal and its negation. Each step in the chain is given an evidence status (such as certain, default, possible). The evidence for both the goal and its negation receives such a status and a resolution procedure then determines the evidentiary status of the goal. In the following example, the rule has a default status, while the fact is certain. The result of combining the default rule with the certain fact is a default status conclusion.

default inference rule: bird(x)→fly(x)
actual fact: bird(tweety)
default conclusion: fly(tweety)

At the same time, ATT-Meta examines evidence for the goal ¬fly(tweety), such as, for example, that Tweety is a penguin or has a broken wing.

Discourse Context

The discourse context consists of beliefs about particular people, places, events, etc. (e.g. author and addressee), about ontological classes of people, objects, events, etc., about language use and

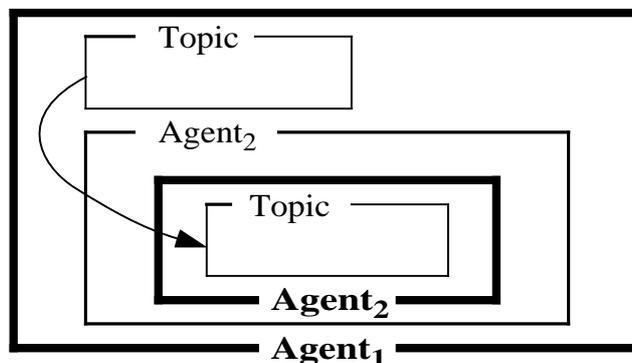


Figure 1: Agent1's beliefs about a Topic and about Agent2's beliefs about the Topic

communication and about social and cultural conventions.

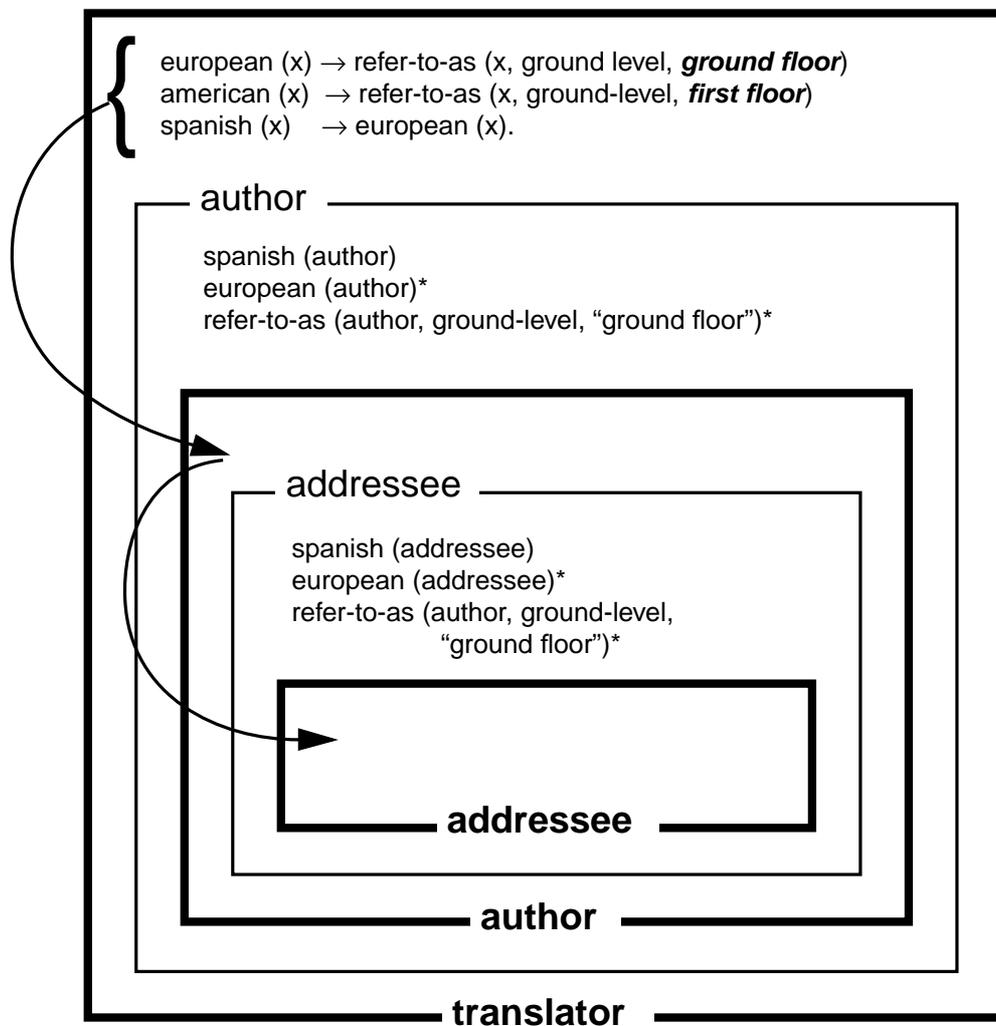
With respect to the specific example under consideration, the discourse context would include such beliefs as:

in Europe (and elsewhere), people refer to the ground level of a multi-story building as the *ground floor*, the next level up as the *first floor*, and so on, in the USA (and elsewhere), people refer to the ground level of a multi-story building as the *first floor*, the next level up as the *second floor*, and so on.

It would also include such beliefs as:

the author is a Spanish speaker, the addressees are Spanish speakers, Spanish speakers may be Spaniards, Spaniards are Europeans,

This knowledge may be accessed by the translator, ascribed to the author or addressees of the source language text or to the audience of the translation in order to establish contextually coherent interpretations or to select contextually relevant expressions. We might represent the contribution of the discourse context for processing the example as in Figure 2 below.



* belief resulting from default inferencing using ascribed rules of inference

Figure 2: Discourse Context

Utterance Context

The utterance context consists of beliefs about the objects and events mentioned or implied during the discourse, the communicative state of the discourse and the “open issues” (i.e., the objects or events whose connections to the context have yet to be established).

With respect to the specific example under consideration, the utterance context would include such beliefs as:

the commercial real estate market in Moscow is expanding rapidly, properties are renting at the equivalent of \$700 to \$800/m²/year, properties are renting at the third highest rates in the world (behind Tokyo and Hong Kong), properties are in high demand, the market is dominated by poverty, legal uncertainty, the principle of the “rich get richer”.

This knowledge may also be accessed by the translator, ascribed to the author or

addressees of the source language text or the audience of the translation in order to establish contextually coherent interpretations or to produce contextually relevant expressions. We might represent the contribution of the utterance context for processing the example as in Figure 3 below.

Interpretation

The process of translation begins with the translator identifying of the intention of the author of the source language text. The translator starts with the assumption that what the translator knows about source language conventions, about cultural and social conventions of source language speakers and about people's presumed knowledge of the world is the same as that which the author knows except where the translator, explicitly or by implication, knows that the author's knowledge differs. Similarly the translator assumes that the author's knowledge about the source

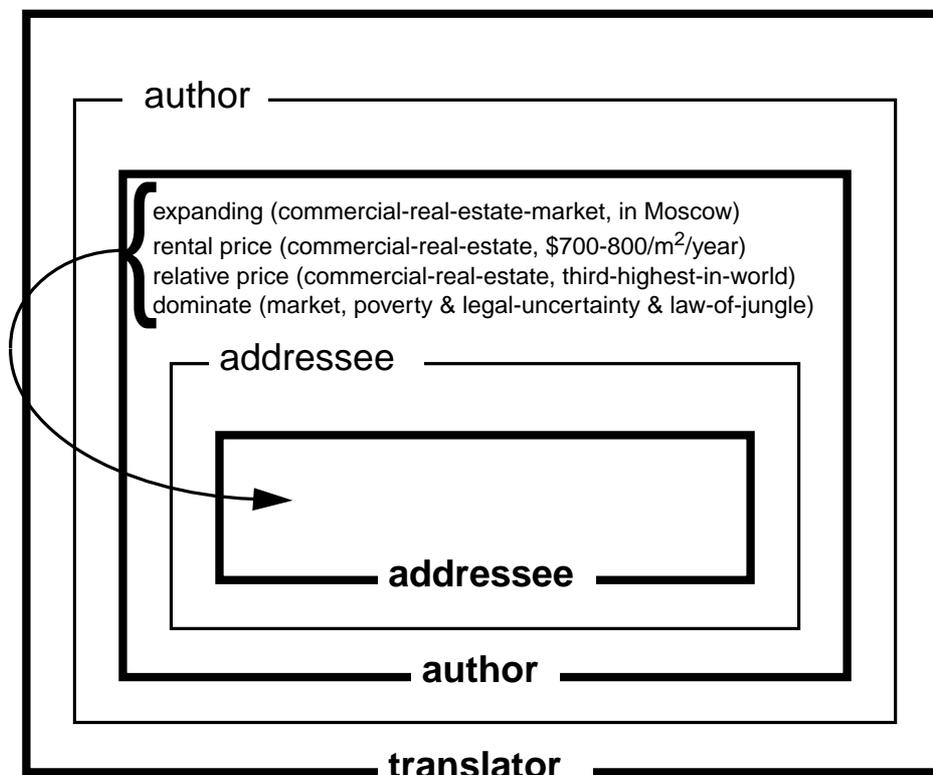


Figure 3: Utterance Context

language addressees is the same as the translator's knowledge of the addressees except where there is explicit or implicational evidence to the contrary. This knowledge, the translator's beliefs modified by that which the translator knows the author believes in contradiction to the translator's beliefs, constitutes the discourse context or background context.

The utterance context provides foreground context for interpretation. The initial context for processing our example might include some of the translator's beliefs about the participants in the source language interaction or about the objects and events described in the text thus far (such as the belief that the Moscow referred to is in Russia, etc.). These are represented below as b_1 , b_2 , and b_3 . It will also include beliefs the translator assumes the author has about the objects and events described in the text thus far (such as the belief that the real estate market in Moscow is booming, etc.). These are represented as b_4 and b_5 .

The initial task of the translator, then, is to assign to the author, for each expression E_i uttered, an intention based on the form and semantics of the expression and the beliefs present in the utterance context. That is, the expression, E_i , is associated with a semantic representation, $p(a,b)$, that needs to be interpreted within the utterance context. For our example, the expressions *el tercer piso* and *el segundo piso* would be provided with semantic representations akin to $\lambda x | \text{floor}(x) \wedge \text{third}(x)$ and $\lambda x | \text{floor}(x) \wedge \text{second}(x)$.

The interpretation involves inferring a belief (represented alternatively as b_6 , b_8 or b_{10}) that is informative and compatible with the utterance context by inferencing from beliefs in the utterance and discourse contexts (represented alternatively as b_7 , b_9 or b_{11}). In other words, b_6 can be inferred from $p(a,b)$ and b_7 as represented in Figure 4, b_8 from $p(a,b)$ and b_9 as represented in Figure 5, and so on. For our example, the translator might infer that the author is referring to the fourth and third levels of the

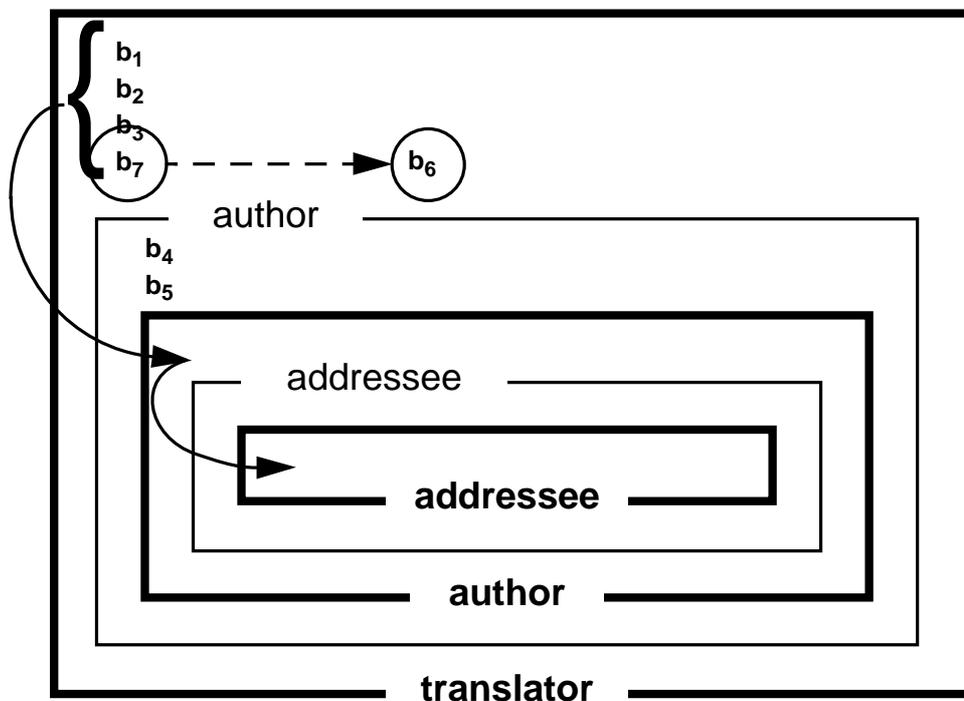


Figure 4: SL interpretation 1 (SL—i1)

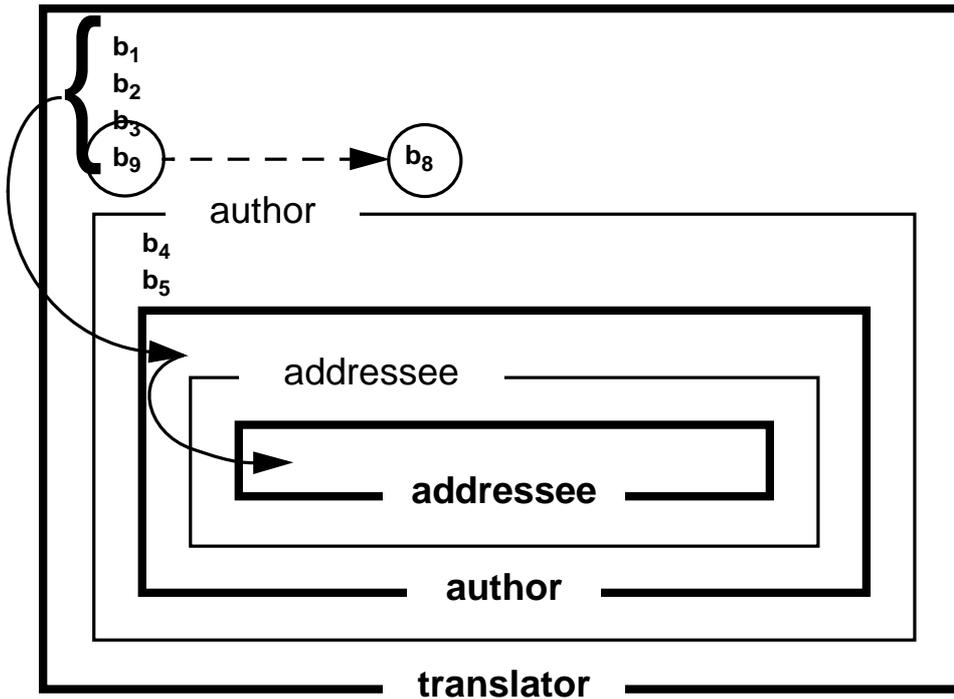


Figure 5: SL interpretation 2 (SL — i2)

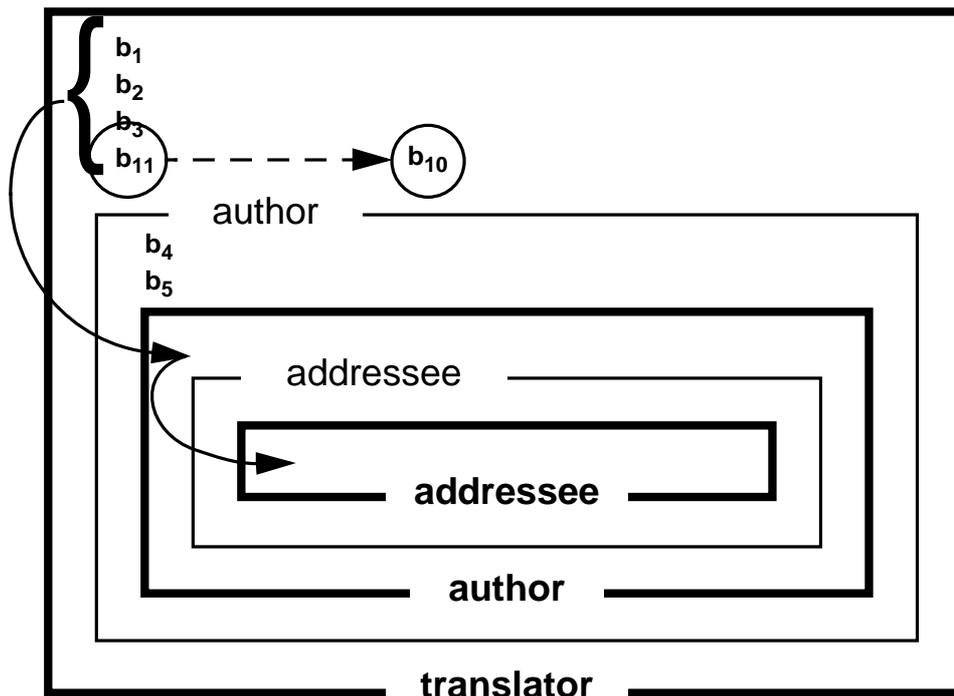


Figure 6: SL interpretation 3 (SL—i3)

building (b₆) given that the author and addressees of the source language text are Spanish speakers (b₂), many Spanish speakers are European (b₃), and Europeans refer to the ground level of a multi-story

building as *the ground floor*, the first level up as *the first floor* and so on (b₇). Other interpretations (e.g., b₈ or b₁₀) are possible given alternative assumptions about who believes what (e.g., b₉ or b₁₁). Finally, from

these possible interpretations, the translator selects one as the author's intention. Assume in this case that interpretation is the one described above (b_6) and represented below as SL—i1.

Translation

The second step is for the translator to express in the target language, with its different set of linguistic conventions, to the TL audience, with its different social and cultural conventions and, perhaps, different conventional knowledge of the world, the SL author's intention (b_6). The translator must produce an expression which permits at least that intention to be identified and, perhaps, any other possible intentions supported by the original SL text.

To begin this process, the SL utterance context prior to the analysis of E_i is subtracted from the author's intention, SL—i1. What remains represents the information added in processing E_i , including the beliefs needed to infer the author's intention from $p(a,b)$.

The next step is to substitute the discourse context of the SL interaction with the discourse context of the TL interaction and to substitute the utterance context of the SL author for the utterance context of the

translator (as TL author). This may result in a change in the beliefs in the TL utterance context and certainly entails changes in the TL discourse context. For instance, suppose that as a result of swapping out the SL discourse context for the TL discourse context, b_7 , the European floor naming convention is blocked in the TL audience's beliefs space by the preexisting b_9 , that is, the American floor naming convention. That is to say, the translator realizes that the floor naming convention of the audience of the translation is different from that of the addressees of the SL text.

The next step is to generate an utterance E_i^* , with semantic representation $p^*(a^*,b^*)$ such that one of its interpretations in the new context is equivalent to the information derived from E_i in the original SL context, i.e., b_6 . Assuming that the intention of the translator (as TL author) is to refer to the fourth and third levels above ground and assuming that the floor naming convention of the TL audience in the American one, the relevant expressions will have semantic representations akin to $\text{ix} \mid \text{floor}(x) \wedge \text{fourth}(x)$ and $\text{ix} \mid \text{floor}(x) \wedge \text{third}(x)$, that is, *the fourth floor* and *the third floor* respectively. This result is represented in Figure 7.

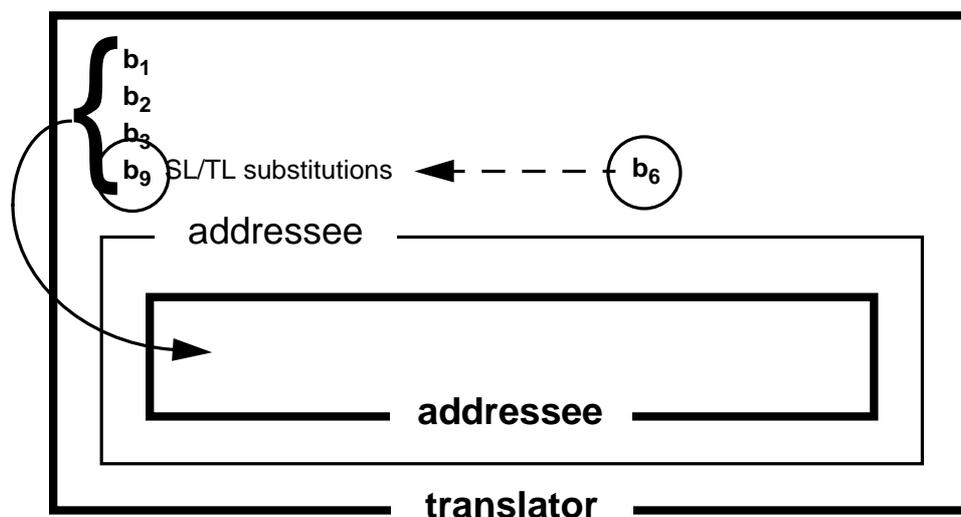


Figure 7: TL interpretation 1 (TL—i1)

The final step is to provide the possible alternative interpretations of E_i^* , given the initial TL discourse context, and TL utterance context.

Alternative Beliefs Affect Readings

There are many potential scenarios beyond those presented above which are based on the observation that the author used Spanish to address the readers of the original article. For instance, with respect to interpretation, the translator might consider as significant the fact that the building referred to is in Moscow. This, coupled with the assumption that Russians appear to follow the US floor-naming convention, implies that the author was referring to the third and second levels above ground in the original text. Alternatively, the translator might consider as significant the fact that the real estate agent that was quoted was speaking Russian. This, again coupled with the assumption that Russian speakers appear to follow the US floor-naming convention, implies that the author (indirectly through the real estate agent) was referring to the third and second levels above ground.

With respect to translation, the translator might take as relevant the fact the text is intended for an audience of real-estate agents and thus the need to identify the exact level above ground is important. In that case the translator might have produced *the third or fourth story* and *the second or third story* as possible translations. Then again, the translator might take as relevant the fact the audience of the translation is uninterested in the precise story referred to but rather in the general availability of space. In that case the translator might appropriately translate the expressions as *one floor* and *the floor below* respectively. We refer to variations in translation which are derived from attempting to accommodate the beliefs of the audience of the translation as "user-friendly" translation.

Thus, those facts identified by the translator as relevant to establishing the author's intention or to expressing that intention to a given audience are crucial to the process. In addition, since most such assumptions are simply not evaluable empirically, it is not obvious that there is any clear notion of "correct" translation.

Note also that an author's intention is not necessarily to describe an event or state of affairs or to present their thoughts on some topic. It might just as well be to communicate a mood or emotion, some sensory input, and so on. Language as a resource for communication provides the author with form as well as meaning which can be manipulated to such ends. Thus, form cannot be discounted a priori in translation in the process.

Finally, a "context", while actually very specific in comparison with all possible contexts for all possible expressions for achieving all possible purposes, can be rather large, abstract, and/or vague (e.g., providing coherence with respect to a Marxist or Freudian or Catholic framework or some combination of such frameworks). That is, if a translator can identify one or another such frameworks as an organizing principle to the written or spoken discourse, especially if it can be expressly attributed to the author, then it too may have to be taken into account.

Equivalence

A central objective of the approach is to support a notion of translation equivalence which is determined by the degree of overlap between information that is stated explicitly and that is inferred from the context in the corresponding source and target language utterances or texts. That is to say, we assume that, because of information gain and information loss during translation, corresponding source and target language utterances or texts are to a greater or lesser degree "equivalent"

depending on the total information conveyed.

The processing schema described above can be used to formally define translation equivalence not on the basis of the meaning of a text, but with respect to the interpretation of the text. We take the interpretation of the source text to be the final context of source addressees' beliefs state less the initial state. This is essentially the elements explicitly communicated by the utterance plus the assumptions and inferences used to support the interpretation.

By comparing the interpretation (core statement plus requisite context) of the source language utterance with the corresponding interpretation of the target language translation, we can determine, beliefs environment by beliefs environment, which beliefs are shared and which beliefs are not shared. Thus, for this simple example case, we end up with an ordered quadruple consisting of (a) beliefs in source addressees' environment that are not in target addressees' environment; (b) beliefs in target addressees' environment that are not in the source addressees' environment; (c) beliefs in both environments with possibly varying levels of confidence; (d) confidence levels for the interpretations as a whole.

This approach allows us to develop a far more elaborate notion of translation equivalence than has been proposed in the past. First, we can distinguish between the core equivalence of expressions in source and target language texts and overall equivalence. Core equivalence is the similarity of speaker's intention for the corresponding interpretations in source and target languages (e.g., SL—i1 and TL—i1 or the levels above ground and the beliefs/inferences required to identify it). The overall equivalence is the total similarity (and difference) of the corresponding sets of interpretations (e.g., SL—i2, SL—i3, ... and TL—i2, TL—i3, ...).

It also allows us to develop both quantitative and qualitative notions of equivalence. Quantitative notions of equivalence take into account the number of beliefs/inferences needed to establish the speaker's intention (or, more broadly, the full set of interpretations). Qualitative notions of equivalence can be defined on the basis of the beliefs/inferences used; their "currency"; the "simplicity" of their connection and so on.

Given this information-rich structure, it is clear that many different kinds of "translation equivalence relations" could be defined. It should also be clear from the foregoing that in very few cases will the interpretations be identical.

Section 3

We have shown in previous sections that beliefs influence translation and that this is reflected in the construction of an interpretation of a text. The interpretation is constructed from the semantics of the text in conjunction with additional propositions needed for integrating that semantics into the utterance context in a coherent manner. It is this interpretation, we claim, which serves as the basis for human translation and which should serve, as well, as the basis for automatic translation.

In this section we examine two further cases which we believe shows that it is just such a coherent interpretation that underlies translation. We look at two different translations of each of two texts. In the first case, patterns of differences in the translations can be related to the different global assumptions of what the story itself is about. In the second case, patterns of differences in the translations can be related to differing world views on the part of the translators.

Background

The texts and translations used for this discussion are drawn from a corpus prepared for the DARPA Machine

Translation Evaluation in 1994, as reported in White *et al.* (1994). These materials are available on-line at:

http://ursula.georgetown.edu/mt_web/index.html

For this evaluation 300 newspaper texts (100 each in Spanish, French, and Japanese) were selected. There was no restriction as to subject domain. Two translations of each text into English were commissioned from two different translation houses. The translators were given strict instructions to neither add nor remove information and to follow a specific style sheet.

The evaluation was conducted entirely by native speakers of English, using one of the English translations as the standard against which the machine translations were evaluated. In addition, the second translation was also evaluated alongside the machine translations.

Translations were rated for fluency, informativeness, and adequacy. Interestingly, when ranked against the first human translation, the second human translation (though better than any of the machine translations) still scored less than perfect in informativeness and adequacy.

We have examined in detail two of the articles drawn from the Spanish corpus along with their accompanying translations. While these articles were selected precisely because they seemed to offer interesting variations in the translations, we do not believe that they differ significantly from the other articles.

Methodology

Our methodology of analysis was as follows. We first processed the Spanish text using the initial modules of CRL's Panglyzer Spanish analysis system (Farwell *et al.*, 1994). This involved first breaking the text into words and sentences and tagging each word and item of punctuation with an identifying part-of-speech tag along

with relevant morphological information. Next, the output is processed by the Panglyzer's phrase recognizer. This groups words in the input into small chunks that are semantically and syntactically cohesive and unambiguous.

Then the two translations were cut up and aligned with these chunks. In some cases where two Spanish chunks were translated by one indivisible English chunk, the two Spanish chunks were combined into one. In other cases, a number of aligned English chunks contained more than one difference between them. These multiple differences were broken out. We also examined the two translations sentence by sentence to look for differences in syntactic structure that were not reflected within any particular chunk.

Finally, we identified and classified each divergent translation unit (and each syntactic difference).

Classification

We had three basic categories of differences: errors, free variation, and belief-based. That is, we felt that any difference in translation either reflected an error or misunderstanding of the text, reflected an arbitrary choice of expression having no impact on the meaning or effect of the text, or reflected a difference in the meaning or effect of the texts due to different interpretations of the source language text. This last category, it should be emphasized, is quite distinct from the error category in that we believe that both of the resulting translations, although different, are licensed by the source language text.

There were also some differences that we refer to as "derivative" in that a difference in one place resulted in a difference in another place. For instance, if one translator used a plural subject where the other used a singular, that could result in a derivative difference in the morphology of the verb.

Note that this classification is based on our understanding of the translator's intent in choosing a particular translation and not simply on the exterior form of the translation. We also note that, as a result, it might be possible in any particular case to argue for a different classification, based on one's own understanding of the text and of the translation. However, the current findings are the result of agreement between the two authors and so have a certain degree of reliability.

Errors. The error category was the smallest of the three classes of differences. This was to be expected since both translations were done by qualified translators. We identified three types of errors. First there are errors that are unintentional or accidental in nature. Spelling errors, for example, generally fall into this category. The second type are errors or extreme awkwardness in the translation due to interference from the source language text. The third class of error consists of intentional errors in that it appears that the translator did not just overlook something but rather got it wrong. That is, the source language text not only did not provide a basis for the translation but rather appeared to provide a basis for a contradictory translation.

Free Variation. The second category we examined was free variation. To some extent this is a flexible category. That is, at the strictest level, there were cases where we could see absolutely no differences in semantic content, connotations, style, register, or invited inferences. These would include such differences as the use (or non-use) of a definite article with plural nouns, writing out numbers versus using numerals, or the use of *that* versus *which* as a relative pronoun. At a less restrictive level, we included cases where there were differences between the lexical choices, but it was not clear that they could be related directly to differing beliefs about the text. An example here is the translation of *sectores* in *sectores costeros* as either *coastal area* or *coastal sector*. *Sector* has a slightly more

military or formal feel than *area*, but not enough for us to classify this as a belief-based difference.

Belief-based Differences. In this category we placed all translation differences that we felt communicated substantially different information, enough so that the readers of the differing translations would have different ideas as to the nature of the source text or the events described in that text. These differences, however, were not such that one could, on the basis of the source text, identify them as wrong or incorrect.

We identified two subtypes of belief-based differences. The first consisted primarily of additions or alterations of information that are related to the beliefs of the translator about the target language audience. That is, information was added or altered if it was felt necessary to communicate the source language author's intent to the target audience properly. Similarly, information may have been deleted if it was redundant and could be recovered from context.

The second subtype of belief-based differences include those that are related to the beliefs of the translators about the events recounted in the source text or on the attitudes of the participants (including the source language author) about these events. Of the two subtypes, these are easily the most common.

Results

The results of our analysis for both texts are presented in Table 1. As can be seen, the results are substantially in agreement. About 40% of the segments in each text showed differences in the translations. Of these differences, there are few outright errors, as one might expect from qualified, professional translators, working under strict guidelines. Of the remaining differences, they are split about half and half between free variation and belief-based differences.

We expected to find a certain number of cases where the different beliefs of the translators about the world (and therefore about the events described in the source text) influenced their translations. And we did. What we did not necessarily expect was that the large majority of these differences patterned on the basis of either differing, yet internally consistent, overall interpretations of the source language article or differing global views of the translators. To demonstrate this, we examine each text in more detail.

Text 1: Earthquakes in Chile

The news event which triggered the original article was an unusual increase in the purchases of living necessities in a Chilean port city, following newspaper and radio reports about the possibility of an earthquake striking the city. The governmental Office of Emergencies (ONEMI) is a major source for this portion (the first half) of the article. The second half of the article is devoted to a more general discussion, supported by quotes from appropriate scientists, about the actual

likelihood of an earthquake and its predictability.

A key factor in forming the translators interpretations of the article appears to be their understanding of the goals and reliability of the Office of Emergency (ONEMI) source. If the goals of ONEMI are understood as helpful, then they may be relied upon as a source of correct information. One may take at face value their criticism of the news media for blowing the earthquake reports out of proportion and encouraging irrational behavior among the populace. The second half of the article validates this view by showing that from a scientific perspective there is both little cause for alarm and not much that one can do anyway in preparation.

On the other hand, if one suspects that the goals of the ONEMI source may be self-serving, then their information may be biased. It could be that the ONEMI source is trying to save face for the organization which currently has egg on it for not having encouraged reasonable preparations for an impending disaster and for not being ready

	Text 1	Text 2
Data		
Number of words	403	392
Number of phrases	170	192
Number of phrases with differences	66	87
Percentage of phrases with differences	39%	45%
Differences		
Errors	13	14
Unintentional		(6)
Interference		(7)
Wrong interpretation		(4)
Free Variation	32	38
Belief-based	36	39
Source Text Related		(30)
Target Audience Related		(6)
Derivatives:		(12)
Total:	81	91

Table 1: Quantitative Results

to deal with the naturally upset citizens. One way to do so is to blame others (in this case, irresponsible media). This view is confirmed by the second half of the article, showing that from a scientific perspective there have been advances in predicting earthquakes and that sensing devices are in place.

These two different evaluations then hook up with different understandings of what kind of story is being written. If one takes the first analysis, then the story fits nicely into a "blame the media" framework in which the accuracy or quality of media reporting of an event itself becomes the subject of another news media story.

On the other hand, if one takes the second analysis, this fits in nicely with a "blame the government" framework in which incompetence by government bureaucrats is the key element of the story, often accompanied by an attempt at covering up the incompetence.

These two broad interpretations of the source language author's intent are encapsulated in the translations of the headline shown below:

Accumulación de víveres por anuncios sísmicos in Chile.

Hoarding Caused by Earthquake Predictions in Chile.

**STOCKPILING OF PROVISIONS
BECAUSE OF PREDICTED EARTHQUAKES IN CHILE.**

In this one phrase there are three differences between the two translations: (1) *Hoarding* versus **STOCKPILING OF PROVISIONS** for *Accumulación de víveres*; (2) *Caused by* versus **BECAUSE OF** for *por* and (3) *Earthquake Predictions* versus **PREDICTED EARTHQUAKES** for *anuncios sísmicos*. In the first translation, irrational, selfish behavior (hoarding) is related by a causal chain (caused by) to an irresponsible speech act (prediction) about a

hypothetical future event. In the second headline, on the other hand, rational activity (stockpiling) is based rationally (because of) on expectations of a future catastrophe (predicted earthquakes). In the first translated headline, it is the predictions of earthquakes that are at the center of the story, while in the second, it is the earthquakes themselves that form the focus for the story.

Another clue to this difference of interpretation is in the translations of:

La posibilidad de un remezón desastroso... tuvo un tratamiento inadecuado en recientes versiones periodísticas,... criticó la repartición.

Here the second translator keeps the structure of the Spanish sentence (grammatically incorrect in English) in which *criticar* (to criticize) takes as its direct object the criticism itself:

The possibility of a disastrous tremor...received inadequate treatment in recent newspaper versions,...criticized the office.

The first translator, however, in keeping with an understanding that the media exaggeration is at the heart of the story, chose to use the verb *say*, but then to indicate explicitly that the report was critical of the media:

The possibility of a disastrous tremor...was inadequately treated in recent news stories,...said the release, which was critical of the media.

The first translator, then, understands there to be an entire report (*repartición*, general sense) issued by the Office for Emergencies as critical of the media, while the second translator envisions only a general comment from the office (*repartición*, usage in Chile), perhaps an attempt on the part of a government bureaucrat to deflect the tough questioning of the reporter.

If these two differing interpretations of the events in the story do, in fact, underlie the two different translations, we would expect to find certain other patterns as well. We would expect, for instance, that the first translator, focussing on the local response to the media misunderstanding, should choose translations which play up the irrational behavior of the local inhabitants, translations that emphasize the unpredictability of earthquakes (and so the implausibility of suddenly preparing for them), and play down the seriousness of such an earthquake, should it occur. The second translator, focussing on the ill-prepared ONEMI, should do just the opposite: play up the rationality of the response of the people to an earthquake that can be accurately predicted to occur soon, with very serious consequences. And, indeed, this is what we find. Without going into all the details, we found that 21 of the 30 source-text related, belief-based differences were supportive of the general hypothesis and its corollaries.

Text 2: Trial in Amazonia

The second text was a news story about the beginning of the trial of a former policeman accused of murdering two brothers, who were instrumental in a labor union in rural Amazonia. The accusation was that the policeman had been hired to commit the killing, and the article goes into some detail about the background of labor violence in the province.

In this article, the results of our study were very similar to those of the first article. There were not many errors and the majority of differences were split between free variants and belief-related differences. The large majority of the belief-related differences (30 out of 36) were related to the source text interpretation, while a smaller number were related to (assumed) beliefs of the target language audience.

As with the first article, we also found that the large majority of the belief-related

differences reflected a broad consistent pattern of differences between the two translations. However, in this case the patterns reflected the two translators' different general world views, particularly those relating to economics and politics. Briefly put, the first translator views the article from what would generally be called a left-wing perspective, while the second translator seems to translate from a more conservative position.

Hovy (Hovy 1988) has shown that the rhetorical goals of speakers influence how their communications of the same event are realized differently in different situations by different speakers and has modeled this for a natural language generator, PAULINE. What we find here is that even translations of a single text describing a particular event differ according to the rhetorical stance of the translator and, in this case, according to the translator's political and economic understanding of the world.

So, for instance, the first (left-wing) translator translates *sindicalista* (describing the murdered brothers) as *labor leader* while the second (right-wing) translator simply uses *union member*. In translating the Spanish *asesino* and *asesinado*, the policeman on trial is described by the first translator as an *assassin* who has *assassinated* the victims, while the second translator describes him as a *killer* who has *murdered* the victims. The first translator translates *campesino* as *peasant* and *terrateniente* as *landholder* while the second translates these terms as *small farmer* and *landowner*. Similarly, the second translator shows more respect by translation *policía* as *police officer*, *justicia* as *justice*, and *jueza* as *the Court* in contrast to the first translator's *policeman*, *law* and *judge*, respectively.

In short, the first translator translates the article from the perspective of a leftist political analysis, where the economically-deprived working classes are in conflict with the upper-class wealthy landowners,

who can use the apparatus of the government to enforce their class interests, and, in the face of strong opposition, may even resort to extra-legal means to maintain their position.

The second translator sees the article from a more conservative perspective. The trial is bringing the killer to justice, despite an atmosphere of violence brought on by the large scale union organizing.

In all, 30 out of the 37 source text-related, beliefs-based differences reflected this underlying difference in world view.

Summary

Overall the important conclusions we draw from these analyses are as follows. (1) Even given strict instructions, competent translators will frequently produce different translations, even of simple newspaper texts. (2) Many of these differences result in translations that are not synonymous. They may differ in actual content, emotional effect, or implied presuppositions. Nonetheless both translations reflect responsible interpretations of the source text. (3) In the two articles examined, most of these meaningful translation differences could be related to an underlying difference in interpretation or understanding of the text as a whole.

Conclusions

With respect to translation, a pragmatics-based approach provides a much more explicit framework for reasoning about the many choices that translators must make in producing a translation. But the central assumption of the approach is that language is vague and texts radically underspecify the interpretation. This is why translators must interpret utterances against a context of beliefs about the world, about the elements of the utterance context, and about the topic and related individuals and states-of-affairs.

The primary implication of this approach to translation is that the beliefs of the translator and the way in which the translator reasons with them account for the eventual form of the translation. Those beliefs and that reasoning should therefore be the focus of any critical analysis of translation quality. We need to ask:

- which beliefs are supported by the text (and how),
- which beliefs are supported by world knowledge (and how),
- which beliefs are supported by linguistic conventions (and how),
- which beliefs are supported by socio-cultural conventions (and how).

If the answers to these questions are satisfactory, the translation is satisfactory. If they are less than satisfactory, then so too is the translation.

From the perspective of a pragmatics-based model of translation and in view of the widespread and significant translation variants to be expected from both human and machine translation systems, it should be clear that the focus of evaluation should be on (1) the similarity and difference between the beliefs of the participants and the inferences performed during the source and target language interactions, and (2) on the naturalness of expression of the target language text. It should also be clear that there is a wide range of potentially appropriate translations for a given interaction. Since variations in translation arise from differences in participants' beliefs and since each of the participants (translator, author, reader and audience) has a different and incomplete knowledge of the individuals, objects, situations and events referred to in a communicative interaction, the potential for variation is quite large.

Finally, we feel that a pragmatics-based approach to machine translation offers the only direct assault on the issues raised by Bar-Hillel as early as 1959 (Bar-Hillel, 1960). It is not simply that MT systems

need knowledge, they need to be able to create complex structures of beliefs and to be able to reason within those structures in order to arrive at an appropriate interpretation in spite of incomplete or possibly inconsistent knowledge. To ignore this fact is to delay progress on both theoretical and applied MT. Form-based translation cannot work: It must be supported by inferencing from knowledge within a context.

References

- Arcand, D. (1989) *Jesus of Montreal*. Miramax: Orion Home Video.
- Attardo, D. (1994) MIKROKOSMOS TMR Builder's Guide. Internal Memo, Computing Research Laboratory, New Mexico State University, Las Cruces, NM.
- Austin, J. L. (1962). *How To Do Things With Words*. Cambridge, MA: Harvard University Press.
- Ballim, A. and Wilks, Y. (1991). *Artificial Believers: The Ascription of Belief*. Hillsdale, NJ: Lawrence Erlbaum.
- Bar-Hillel, Y. (1960). The Present Status of Automatic Translation of Languages. In F.L. Alt (ed.) *Advances in Computers*, 1:91-163. New York: Academic Press.
- Barnden, J.A., Helmreich, S., Iverson, E. and Stein, G.C. (1994). An Integrated Implementation of Simulative, Uncertain and Metaphorical Reasoning about Mental States. In J. Doyle, E. Sandewall and P. Torasso (Eds.), *Principles of Knowledge Representation and Reasoning: Proceedings of the Fourth International Conference*, pp.27-38. San Mateo, CA: Morgan Kaufmann.
- Carlson, L. and Nirenburg, S. (1990). World Modeling for NLP. *Technical Report 90-121*. Center for Machine Translation, Carnegie Mellon University, Pittsburgh, PA.
- Farwell, D., and Helmreich, S. (1993). Insights into Communicative Intent from Pragmatics-based Translation Equivalence. Paper presented at *The Seventh Annual Southwestern Cognitive Science Conference*. March 6-7, 1993. Tucson, AZ.
- Farwell, D., and Helmreich, S. (1995). This is not a Bedroom Farce: Pragmatics and Translation. *Proceedings of the Fourth International Colloquium on Cognitive Science (ICCS-95)*, pp.73-82. Donostia, Spain.
- Farwell, D., and Helmreich, S. (1996). Ground Floor? First Floor? Going Up. Presented at the *Fifth International Pragmatics Conference*. July 4-9, 1996. Mexico City, Mexico.
- Farwell, D., and Helmreich, S. (1997). Assassins or Murderers: Translation of Politically Sensitive Materials. Presented at the *26th Annual Meeting of the Linguistics Association of the Southwest International Pragmatics Conference*. October 4-9, 1997. Los Angeles, CA.
- Farwell, D. and Helmreich, S. (1998). Translation Differences and Pragmatics-Based Machine Translation. *Machine Translation* 13: 17-39.
- Farwell, D., Helmreich, S., Jin, W., Casper, M., Hargrave, J., Molina-Salgado, H., and Weng, F. (1994). Panglyzer: Spanish Language Analysis System. *Proceedings of the First Conference of the Association for Machine Translation in the Americas*, pp.56-65.
- Fauconnier, G. (1985). *Mental Spaces: Aspects of Meaning Construction in Natural Language*. Cambridge, MA: MIT Press.
- Fodor, J. A. (1984). *The Modularity of Mind*. Cambridge, MA: MIT Press.
- Green, G. M. (1984) Some remarks on how words mean. Indiana University Linguistics Club, Indiana University, Bloomington, IN.
- Grice, H. (1975). Logic and conversation. In P. Cole and J. Morgan, (Eds.) *Syntax and semantics*, 3: Speech acts, pp. 41-58. New York: Academic Press.

- Hovy, E. (1988). *Generating Natural Language Under Pragmatic Constraints*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Lakoff, G. (1987). *Women, fire, and dangerous things*. Chicago: University of Chicago Press.
- Lakoff, G. and Johnson, M. (1980). *Metaphors we live by*. Chicago: University of Chicago Press.
- Nirenburg, S. (Ed.) (1987). *Machine translation: Theoretical and methodological issues. Studies in Natural Language Processing*. Cambridge: Cambridge University Press.
- Nirenburg, S., Carbonell, J., Tomita, M., and Goodman, K. (1992). *Machine Translation: A Knowledge-Based Approach*. San Mateo, CA: Morgan Kaufmann.
- Nunberg, G. 1978. *The Pragmatics of Reference*. Indiana University Linguistics Club, Indiana University, Bloomington, IN.
- Onyshkevych, B. and Nirenburg, S. (1994). The Lexicon in the Scheme of KBMT Things. *Memoranda in Cognitive and Computer Science: MCCS 94-277*, Computing Research Laboratory, New Mexico State University, Las Cruces, NM.
- The PANGLOSS Mark III Machine Translation System (1994). A Joint Technical Report by NMSU CRL, USC ISI and CMU CMT. Carnegie Mellon University, Pittsburgh, PA.
- Pustejovsky, J. (1991). The Generative Lexicon. *Computational Linguistics*, 17:4.
- White, J., O'Connell, T., and O'Mara, F. (1994). The ARPA MT Evaluation Methodologies: Evolution, Lessons, and Future Approaches. *Proceedings of the First Conference of the Association for Machine Translation in the Americas*, pp.193-206, Columbia, MD.
- Wilks, Y. (1975). A Preferential Pattern-Seeking Semantics for Natural Language Inference. *Artificial Intelligence*, 6:53-74.