



ELSEVIER

Journal of Pragmatics 30 (1998) 259–274

journal of
PRAGMATICS

On the nature of blending as a cognitive phenomenon[☆]

Francisco José Ruiz de Mendoza Ibáñez*

*Departamento de Filologías Modernas, Universidad de La Rioja, c/Cigüena 60,
26004 Logroño, La Rioja, Spain*

Received 26 January 1996; revised version 26 May 1997

Abstract

This paper argues for the integration of Turner and Fauconnier's concept of blending within the general architecture of cognition provided by production models. It also revises the notion of blending as proposed by Turner and Fauconnier and observes that blending obeys general constraints of cognitive economy and of consistency; more specifically, it is argued that blending responds to Lakoff's Invariance Principle, of which a slightly modified version is provided here and defined as the Extended Invariance principle. The Extended Invariance principle accounts for the possibility of preservation of all generic-level structure (not only image-schematic structure) in metaphoric mappings, and allows for the possibility of more than one input space in the mapping process. Finally, blended spaces are seen as one possible outcome of the activity of working memory, while the apparent conceptual irregularities and asymmetries which Turner and Fauconnier observe in their emergence are alternatively explained as contextual implications motivated by and derived from the convergence of information from different Idealized Cognitive Models (or ICMs) in the action part of productions.

1. Introduction

Mark Turner and Gilles Fauconnier have recently proposed blending as a cognitive process that operates over mental spaces and whose operational value ranges over probably all cognitive phenomena (see, for example, Fauconnier and Turner, 1994, 1996; Turner and Fauconnier, 1995a,b). This includes categorization, inference, metaphor, analogy, and the making of hypotheses among others. Blending consists of the conceptual projection of two mental spaces (the source and target) into a third, 'middle' space called the 'blended' space. A mental space is "a small conceptual packet constructed as we think and talk" (Turner and Fauconnier, 1995: 30). A

[☆] Financial support for this research has been given by the DGICYT, grant no. PB92-0757, Ministry of Education and Science, Spain. I am grateful to two anonymous referees and to Jacob Mey for their very useful comments.

* E-mail: franruiz@dfm.unirioja.es

blended space typically contains specifics from source and target which, according to Turner and Fauconnier, also include non-correspondences which yield an impression of counterfactual or impossible structure. In addition to the three spaces mentioned, there is one further middle space called the ‘generic’ space. This contains abstract information consisting of the roles, frames and schemas common to the source and target spaces. This information is then available for projection to other possible targets. The authors give proverbs as an example of generic space (see the discussion on the *GENERIC IS SPECIFIC* metaphor in the next section below).

The detailed analysis of different blends has led Turner and Fauconnier to observe some relevant features: blends exploit counterpart connections between input spaces; they develop structure not provided by the inputs; they use a large amount of conceptual structure, though we are often unaware of it, and they can be inputs to other blends. However, it remains to be explained how the projections are carried out, and what principles, if any, regulate the processes involved. In this paper, it is argued that projections are subject to a number of constraints. More specifically it will be claimed that inferences within a blended space are consistent with ‘production structure’ (this term is constructed from the notion of ‘production’ as used in Anderson, 1983, 1987), and that blending processes need to respect Lakoff’s well-known Invariance Principle (see Lakoff, 1990, 1993), of which a slightly modified version will be provided below.

2. Invariance

Lakoff and his collaborators (e.g. Lakoff and Johnson, 1980; Lakoff and Turner, 1989; Lakoff, 1993) have developed a theory of metaphor in which metaphor is seen as a systematic conceptual mapping from a source domain to a target domain. One well-known example he gives is the *LOVE IS A JOURNEY* mapping, which consists of the following correspondences (Lakoff, 1993: 207):

- The lovers correspond to travellers
- The lovers’ common goals are their common destinations
- The love relationship is the vehicle
- Difficulties in the relationship are impediments to travel

These correspondences explain every-day English expressions like *Look how far we’ve come*, *It’s been a long, bumpy road*, *Our relationship is off the track*, and others. A journey is the source domain that helps us to talk about love (the target domain) by means of the mapping.

The source domain is sometimes characterised as an image-schema. We have image-schemas of a bounded space (container), of a path, of contact, and of human orientations (up–down, front–back, centre–periphery) (Lakoff and Turner, 1989: 97). These we apply to our understanding of certain situations. For example, the image-schema of a path explains our understanding of death as departure (e.g. *He passed away*, *the departed*); being out (of the centre of functional interaction) is

associated with non-functionality (as in *be out of order*, *be spaced out*, *be worn/zonked out*, etc.) in the centre-periphery orientation; emotional states respond to the container schema (e.g. *be in love*, *fall out of love*); the metaphor MORE IS UP (which is the basis of expressions like *high prices*, *high speed*, *high number*, etc.) is based on the up-down schema (see Lakoff, 1987: 271ff.; Taylor, 1989: 134ff.).

Lakoff (1990, 1993) has argued that when metaphors map spatial domains onto non-spatial domains, the structure and logic of the image-schemas involved needs to be preserved. This requirement has been termed the Invariance Principle (Lakoff, 1993: 215):

The Invariance Principle:

Metaphorical mappings preserve the cognitive topology (that is, the image-schema structure) of the source domain, in a way consistent with the inherent structure of the target domain.

This means that, for example, for a container-schema a source domain interior is only mapped onto a target domain interior, never to an exterior; for a path-schema, a goal will map onto a goal, never onto a source or a trajectory, and so on. If you take the LOVE IS A JOURNEY metaphor you will see that in the correspondences, destinations become goals, not vehicles or lovers. The Invariance Principle sets constraints on correspondences.

The importance of the Invariance Principle can hardly be overstated. Lakoff (1990) has been able to provide suggestive evidence that image-schemas are an essential part of conceptualization. For example, we understand such abstract concepts as time, states, changes, causes, purposes, scales, and categories in terms of image-schemas.¹ Lakoff also provides examples of the application of the Invariance Principle in the cases of image metaphors and of generic-level structure. In the former, the interpretation seems to require the previous analysis of the image in terms of the related image-schema. Consider the following example (Lakoff, 1993: 229):

(1) My wife ... whose waist is an hourglass

A woman and an hourglass have much of their conventional overall shape in common. This allows the metaphor to map the middle of an hourglass onto the waist of a woman; parts may be mapped onto parts provided that they stand in a similar structural relation to their respective wholes. We may add that no part of the hourglass may be mapped onto the wife's limbs since this would violate the restrictions imposed by the Invariance Principle on the nature of correspondences.

¹ Time, states, and categories are often seen in terms of the container schema. Since change is seen as motion to a new location, causes as forces, and purposes as destinations, these concepts invoke the path-schema. Finally, scales are interpreted in connection to the up-down orientational schema. Other alternative conceptualizations are possible: for example, the path schema is the basis for our understanding of time as a moving entity (cf. *in time* versus *Time went away*), and of states as possessions (cf. *I'm in trouble* versus *I have trouble*). For a study of the role of image-schemas in characterizing event structure, see Lakoff (1990: 57–65).

The analysis of generic-level structure also proves interesting. Lakoff and Turner (1989) have studied how generic-level mappings like personifications and proverbs preserve image-schematic structure. Personifications seem to fit a pattern where events are understood in terms of actions performed by a personified agent. Thus, Lakoff and Turner postulate the existence of the EVENTS ARE ACTIONS mapping, which is constrained in such a way that the causal structure, the aspectual structure, and the persistence of entities are preserved. For example, the event of death may be understood in terms of the action of reaping. In the resulting metaphor death is personified as a reaper. Proverbs work similarly. Consider Lakoff and Turner's (1989) analysis of the GENERIC IS SPECIFIC metaphor. This maps a single specific-level schema onto other specific-level schemas that share the same generic-level structure. Thus, this metaphor allows us to understand a number of situations in terms of one particular situation. For example, in *Blind blames the ditch*, the blind person represents any person with some sort of incapacity (physical, mental, etc.), the ditch is the situation that the person is unable to cope with because of his incapacity, and the act of blaming represents the foolish attitude of holding the situation responsible for failure rather than the person's inability. The generic structure extracted from the proverb maps onto the structure drawn from any of its potential applications, and in so doing preserves the same event shape.

The Invariance Principle deals exclusively with image-schematic structure or with generic-level structure interpreted in terms of image-schematic structure. However, in its present formulation, it falls short of providing a satisfactory explanation of other consistency phenomena existing between source and target domains. Consider this example, drawn from Blakemore (1992: 163):

(2) My neighbour is a dragon

A woman is usually called a dragon when she is fierce and unpleasant. There are several points which need to be made about this metaphor. First, as observed by Blakemore, by using the metaphor rather than a paraphrase like *My neighbour is fierce and unfriendly*, the speaker puts the hearer to some extra processing effort which is only justified if it is compensated by some extra contextual implications which yield some extra contextual effects.² In the metaphor, these implications will be about the nature of the neighbour's behaviour and her appearance. Second, as observed by Ruiz de Mendoza (1997), the mapping in this sort of metaphor (which belongs to the PEOPLE ARE ANIMALS system) differs substantially from the mapping in metaphors of the LOVE IS A JOURNEY kind. In LOVE IS A JOURNEY, the source domain is a concrete domain of experience whereas the target is abstract; in PEOPLE ARE ANIMALS, in contrast, both domains are concrete. As a result, the mapping in the former

² Blakemore follows Relevance Theory, as expounded by Sperber and Wilson (1986). In it, human communication and cognition are seen as governed by the search for relevance. Relevance is achieved as the result of a balance between contextual effects and processing effort, in such a way that the greater the contextual effects, the greater the relevance, but the greater the processing effort to obtain these effects, the lower the relevance. Contextual effects are the result of the interaction of newly-presented information with a context of existing assumptions. This may happen by strengthening, by contradicting, or by combining with an existing assumption. In the last case, we have a contextual implication.

serves a structuring function of the whole target domain, while in the latter the mapping is used to single out and place in focus one aspect of the target which is, by convention, particularly relevant; the rest of the domain remains unaffected. Third, there is generic-level structure involved in understanding this metaphor. Thus, Lakoff and Turner have accounted for the PEOPLE-ARE-ANIMALS metaphor in terms of a more generic one, the GREAT CHAIN metaphor. The Great Chain is a cultural model which defines attributes and behaviour applying to humans, animals, plants, complex objects, and natural physical things (Lakoff and Turner, 1989: 170–171). By linking the Great Chain to the GENERIC-IS-SPECIFIC metaphor – which picks out from specific schemas common generic structure – we are able to understand general human character traits in terms of well-understood non-human attributes and, conversely, to understand some non-human attributes in terms of human attributes. Thus, Lakoff and Turner point out, in the metaphor *Achilles is a lion*, the character of Achilles is understood in terms of an instinctive trait of lions, i.e. their conventional quintessential property of being brave. This property is actually the result of another previous metaphor: understanding the behaviour of the lion in terms of the courageous behaviour of a human. The two metaphors cancel each other out but, interestingly, preserve the structure of the two schemas involved: the relation between the lion and his courage is mapped onto the relation between Achilles and his courage. This account also applies to Blake-more's example and explains how its contextual implications are obtained. The relation between the dragon and its attributed behaviour and appearance is mapped onto the relation between the neighbour and her observed behaviour and appearance.

But what is particularly revealing about metaphor (2) – which makes it altogether different from others like LOVE-IS-A-JOURNEY or like (1), for that matter – is the fact that the constraints on the mapping are not image-schematic. Thus, although one could think that the part–whole image-schema is likely to play some role in the interpretation of a metaphor where two physical entities are set in correspondence – as in (1) above –, in (2) only a dragon's attributed behaviour is made to correspond with a person's behaviour. As a result, even though generic-level structure does seem to guide the understanding of the metaphor, as we noted above, this generic structure is neither intrinsically image-schematic nor interpretable in image-schematic terms. The structure is simply propositional. This places an important limitation on the purported broad scope of application of the Invariance Principle. This limitation may be sorted out, however, through a reformulation of the principle, of which a first tentative version – called the Extended Invariance Principle – is provided here:

The Extended Invariance Principle:

Metaphorical mappings preserve the generic-level structure of the source domain in a way consistent with the inherent structure of the target domain.

3. Convergence of cognitive models

To the Extended Invariance Principle we have to add a corollary. Sometimes for full metaphor interpretation we need to invoke knowledge from more than one

knowledge domain (or Idealized Cognitive Model, in Lakoff's terminology).³ In such cases the cognitive structure of these additional domains needs to be preserved too. Let us contrast the metaphors involved in these two utterances:⁴

- (3) (I'm afraid) I'm getting ahead of myself.
 (4) I'm (luckily) ahead of schedule

Ahead of invokes a path-schema which is the basis for the unfolding of the situations in (3) and (4). It describes the relative positions of two sets of entities along a path in a front-back orientation, as in *There were three people ahead of me in the queue*. It also conveys an idea of movement in a same direction (in contrast to *in front of*).⁵ It can be seen that both (3) and (4) tend to describe a situation where the speaker has made more progress than expected in relation to some expected goal. But the situation is felt to be negative in (3) and positive in (4). This is a crucial difference between the two utterances which, as we shall contend, can be only explained by looking into the nature of the conceptual domains involved in the metaphoric mappings.

First, we shall argue that in (3) we have a metaphor belonging to the general system of metaphors which Lakoff (1996) has called THE DIVIDED PERSON. This is the set of general correspondences in Lakoff's proposal:

- A person is an ensemble (the subject plus a self).
- The experiencing consciousness is the subject.
- The bodily and functional aspects of a person constitute a self.
- The relationship between subject and self is spatial.

The metaphors in this system are various, like the loss of self, the split self, the true self, the absent subject, the scattered self, and others. These metaphorical subsystems deal primarily with three issues: control, compatibility, and values. In (3) we have a case of the ABSENT SUBJECT metaphor. This is a metaphorical model used to conceptualize normal self-control by the subject or the lack of it. Normal self-control is seen as keeping the subject within a bounded region or in an upright position (vertical orientation). As a consequence, expressions like *beside myself*, *spaced out*, *floating off*, indicate loss of self-control. It can be observed that each of these expressions exploits a different image-schema (path, container, up-down). Since being ahead of oneself invokes a path-schema and indicates lack of self-control – which in itself is considered to be negative: if we cannot control somebody else's activities we run the

³ Idealized Cognitive Models (or ICMs) have been proposed by Lakoff (1987) to organize our knowledge. More specifically, he suggests that ICMs make use of four possible kinds of structuring principles: propositional structure (e.g. frames), image-schematic structure, metaphoric mappings, and metonymic mappings. For him, category structures and prototype effects are the by-product of this organization.

⁴ The examples have been borrowed from Fauconnier and Turner (1994), but our treatment is different.

⁵ Here we are talking about core meanings. It is possible to use both *ahead of* and *in front of* with our without movement (e.g. *John was walking in front of me*, *Peter stood ahead of me in the queue*). However, in predications where there is no movement predicate involved that overrides the core meanings, *ahead of* implies movement of at least one of the sets of entities involved, while *in front of* preferably describes a static situation (cf. *John was in front of me*, *Peter was ahead of me in the queue*).

risk of experiencing non-beneficial consequences – we can list this metaphor with the others in the subsystem.

For metaphor (4), which is – in contrast – a case of self-control, and in which scheduled events in time are seen as points along a path (TIME-IS-SPACE), we need to propose the complementary activation of what we can call the Control ICM, a propositional model of which we provide a partial description below:

The Control ICM:

- (a) A person controls an entity or a set of entities when it is within the person's power to decide on the way the entity or the set of entities will behave.
- (b) A person controls a state of affairs when it is within the person's power to decide whether the state of affairs will obtain.
- (c) Control generally decreases in proportion to physical distance.
- (d) Maximum control of an entity, a set of entities, or a state of affairs is desirable/ Minimum control is undesirable.

Note that part (d) of the ICM is stated as a scale. This scale interacts with the path-schema in the sense that the greater the distance between the self and the subject, the less control the latter will have over the former. The less control, the more undesirable the situation will be for the subject. Conversely, the greater the control, the more desirable the situation will be.

Consequently, this ICM plays a crucial role in the interpretation of the state of affairs pictured by metaphors (3) and (4). But note that the application of the ICM follows the metaphorical mapping; it is in no way previous to it. Working otherwise would be inconsistent with the conceptual structure of the source domain: if there is no divided self, there is no lack of control; if there is no lack of control there is no undesirable situation. The interaction between the control ICM and the source domain in the DIVIDED PERSON metaphor produces a range of contextual effects without violating image-schematic (i.e. generic-level) structure. In the same way, the image-schematic structure of the Control ICM, which becomes another input domain for the metaphor, is also preserved: greater distance correlates with greater loss of control, and therefore with greater undesirability. The nature of the contextual effects motivated by a metaphoric mapping has to be in accordance with the Extended Invariance Principle, which we are now in a position to reformulate:

The Extended Invariance Principle:

All contextual effects motivated by a metaphoric mapping will preserve the generic-level structure of the source domain and of any other input space involved, in a way consistent with the inherent structure of the target domain.

4. Productions

Our next hypothesis is about the way different input spaces are called to interact in our minds for metaphor production and interpretation in an efficient way. We

shall contend that the interactions are consistent with the requirements of what in cognitive psychology is known as production systems (Anderson, 1983, 1987). In this view, productions provide a blueprint for the convergence of different mental spaces, thereby constraining the nature of this cognitive process. In fact, as will become apparent later, the kind of contextual effects obtained by the convergence of mental spaces and the amount and kind of material to be included in a mental space is probably a function of the nature of the production system under operation.

The nature of productions is to be understood in light of the by now well-established division cognitive psychologists have made between declarative and procedural knowledge. Declarative knowledge is encyclopedic knowledge about the world ('knowing what'); procedural knowledge is the knowledge required to do things, like drive a car, swim, or tie one's shoes ('knowing how'). A production system is a set of conditional-action (or 'if-then') pairs, where the conditional part specifies some condition which is to be met for the action part to be executed. Therefore, production systems are procedural knowledge systems. In order to see the form productions take, consider the following example simplified from Anderson (1983: 264).⁶ This is intended to explain how the phrase *The rich doctor* is generated:

- (5) IF the goal is to describe the object
 and the object is known to the listener
 and the object has a certain property that is to be described
 THEN set as subgoals
 1. to generate *the*
 2. to generate the name for the property
 3. to generate the name for the object

Obviously, the names for the property and the object are retrieved from declarative memory. This is done via working memory which is a kind of short-term memory, central to most thought processes. In Anderson's proposal, two processes relate working memory to declarative memory: *storage* and *retrieval*; other two relate working memory to production memory: *match* (when data in the working memory are put in correspondence with the conditions in the productions), and *execution* (when the matched productions are deposited in working memory). Then, the relationship between working memory and the outside world is based on *encoding* (depositing information about the outside world into working memory) and *performances* (which convert commands in working memory into behaviour).

Anderson's model meets the production and interpretation requirements for metaphors (3) and (4) above. Consider the following tentative productions:

⁶ Production systems are computer programs which model some cognitive phenomena. The systems are usually applied to problem solving and learning skills. Anderson's proposal (see also Newell, 1989, for an alternative view) provides an integrated model of knowledge types which, we contend, is sensitive to the requirements of a theory of mental spaces.

- (6) IF the goal is to describe a situation where
- a person has completed more scheduled activities than expected, and
 - this is felt by the speaker to have negative consequences for the person,
- THEN set as a subgoal
- to generate a metaphorical mapping where
 - each scheduled activity is seen as a point along a path, and
 - the person is seen at two different points along that path.
- (7) IF the goal is to describe a situation where
- a person has completed more scheduled activities than expected, and
 - this is felt by the speaker to have positive consequences for the person,
- THEN set as a subgoal
- to generate a metaphorical mapping where
 - each scheduled activity is seen as a point along a path, and
 - the person is seen at a point along that path closer to the end of the path than expected.

The actions in these productions are based only on a subset of the general set of correspondences which constitute the A CAREER IS A JOURNEY metaphor for (6) and (7), and the DIVIDED PERSON metaphor for (7). It will be noted that cognitive economy is at work in bringing to bear only a limited relevant portion of the full set of correspondences that could be activated and which make up the whole metaphorical ICM. Also, in being part of a production this subset has become procedural knowledge, which permits its quick, effective application.

Then, we have the problem of the negative or positive overtones of the metaphors in question. These were treated above as contextual effects which have their origin in the activation of the Control ICM by implication. Because of their inferential nature, for a case like (6), they may take roughly this form: if the self is at a different point moving away from the subject, then the subject is losing control over the self, which is undesirable (negative overtone). For (7), in contrast, the positive feeling is derived from the following line of reasoning: if a person achieves a goal, then the person gains intended control over a state of affairs, which is desirable; the closer to the goal, the closer the person is to gaining control. It may be observed that inferences share with productions their if–then structure, but they are not productions since they do not involve actions.⁷

A variation in part of the conditions of a production may either call for the activation of a new altogether different production or for a complementary one. For example, for:

⁷ Other inferences are also possible, the only constraint on their production being consistency with the Extended Invariance Principle. For example, metaphor (3) may suggest that the speaker feels he is not doing things at the right time. Remember that we have a situation where it is the self that is not doing things according to schedule. The subject knows the right time and would like to abide by it. Consequently, only if the subject is capable of keeping the self under control, will the appointed schedule be observed.

(8) He's getting a little bit ahead of himself on his way to promotion,

where the nature of the path is specified, we need a complementary production:

- (9) IF the goal is to describe the nature and aim of the scheduled activities
 THEN set as a subgoal
 to generate a metaphorical mapping where
 – the nature of the activities is seen as a type of path, and
 – the aim of the activities is seen as the end of the path.

The idea here is that any system will apply by default the more general production if that is felt to be sufficient for the purpose in hand. If more specificity is needed this will require further activations. Note also that, as should be expected, the correspondences used in the productions obey the Extended Invariance Principle: that is, path types map onto activity types, and destinations map onto goals.

Finally, it is also worth noting that if the addressee seeks to explore further the contextual effects of a mapping, he will need to correlate each of the actions in the THEN-part of the production with the appropriate ICM which is to be retrieved from declarative memory (in our example part of the Control ICM, which will become a complementary input space). But there is no apparent reason why one should need to make use of any ICMs to be retrieved from declarative memory as long as production and interpretation can be performed on the basis of production systems. Again, this would agree with the general economy constraint which seems to pervade cognitive activity.

5. Blending

According to Turner and Fauconnier (1996: 113), mental spaces are built up provisionally and they undergo revision as the discourse unfolds. In this they match Anderson's view of working memory, which makes it possible to integrate the theory of mental spaces into the general architecture of cognition provided by production memory models. Thus, if we think of declarative memory as containing a large number of ICMs, of procedural memory as containing conditional-action pairs (where the actions are based on selected portions of ICMs), and of working memory as provisionally creating different mental spaces (or provisional knowledge packets) which are acted upon in production and interpretation, we shall be able to shed some light on some of the cognitive processes associated with metaphorical mappings.

For example, consider the matching thought process described above. This makes sure that the data in working memory, which are derived in part from the outside world and in part from declarative memory, are consistent with the conditions in the productions. If there is any inconsistency between these data and the conditions, the production will have to be abandoned and we shall have to look for a new correct one. If no other matching production is found, we shall be compelled to construct a

new production or to alter the features of a preexisting one. This is probably what happens during the learning process.

Now, think of the action part of productions like (6) and (7) above. These are built from information selected from relevant ICMs. Each information packet thus formed is a mental space, in terms of Turner and Fauconnier's theory. Different mental spaces will then become input spaces for the principle-regulated activity of working memory. As a result, new structures (such as blends) will emerge containing information from the relevant input spaces. However, since the activity of working memory is simply one of selecting and combining information from other long-term knowledge stores, no new structure is likely to violate general constraints like the Extended Invariance Principle. If this hypothesis is correct we need to revise part of Turner and Fauconnier's theory since they contend that the Blend creates its own structure which may be inconsistent with the structure of some of the input spaces. Let us consider Turner and Fauconnier's (1995b) discussion of the metaphor

(10) You're digging your own grave

which they use, among other examples, in their argument to claim that blended spaces may create structure inconsistent with the information in the input domains.⁸ They observe that this is a conventional expression used as a warning. It suggests that the addressee's behaviour is going to have disastrous consequences and that he is unaware of this. Failing is seen as being dead and buried; bad moves that precede and cause failure correspond to grave-digging. Also, just as it is foolish to facilitate one's own burial by digging one's own grave, it is foolish not to be aware of the consequences of one's own actions. The metaphor is, therefore, based on the mapping from the domain of graves, corpses and burial to the domain of getting into trouble, being unaware of one's actions and failure. But, according to Turner and Fauconnier, there are important mismatches between the source and the target of this metaphor:

- The causal structure is inverted: foolish behaviour causes failure but grave-digging is not the cause of death.
- The intentional structure does not carry over: people who dig a grave are aware of what they are doing, but the figurative digging is an unintentional misconstrual of action.
- The frame structure of agents, patients and sequence of events is not preserved. In real life, the patient dies, the agent digs the grave and buries the patient, and there is no temporal connection between finishing the digging and dying. But in the metaphor, agent and patient are fused and there is a temporal connection between finishing the digging and dying.

⁸ Lack of space makes it impossible to discuss all of Turner and Fauconnier's examples here. Metaphor (11) is one major example found in Turner and Fauconnier (1995b). Another one, the Grim Reaper, is discussed in detail in Ruiz de Mendoza (1996). It may be observed in Turner and Fauconnier's examples that the information and structure which they attribute exclusively to the blend can be easily derived from the long-term knowledge store. That this is the case will become evident as we analyse example (11).

- The internal event structure does not match: in the target the more trouble you are in, the more you risk failure, but in the source there is no correlation between the depth of somebody's grave and his chances of dying.

Turner and Fauconnier explain these apparent mismatches as an idiosyncrasy of the blend. According to them, the blend inherits the concrete structure of graves, digging, and burial from the source, and the causal, intentional, and internal event structure from the target. Both source and target are envisaged as input spaces to the blend, which is seen as having its own specific structure. For example, in the blend digging one's own grave is a serious mistake since it makes dying more probable, and it becomes possible, as in the target but not in the source, to be unaware of one's own actions; also from the source we obtain the idea that it is highly foolish to be unaware of such concrete actions.

But there are some problems with this vision of blended spaces. It is true that in production and interpretation we make use of information from more than one knowledge domain, but we have seen before that the treatment of such information needs to obey the Extended Invariance Principle, i.e., to preserve the generic-level structure of the source domain(s) in a way consistent with the inherent structure of the target domain. If source and target domains abide by this principle, why should our cognitive mechanisms have a special space where all types of irregularities and inconsistencies may come about? Here we shall contend (i) that blends are middle spaces where information is acted upon, as in Anderson's proposal of a working memory; (ii) that the information derives from different folk models which are projected into the middle space and blended there without violating the general principles of consistency between domains.

To defend this view we need to reexamine Turner and Fauconnier's claims about metaphor (10). Let us set it in contrast to other related expressions:

- (11) You're digging your own political grave
- (12) You're committing political suicide

According to Turner and Fauconnier the metaphor in (10) is a warning that carries two implications: the addressee is acting wrongly, which will cause him to have a bad experience, and he is unaware of this, which is foolish. We can assume that (11) would be explained by the authors as a specification of the kind of failure the addressee is expected to experience. Then, (12) conveys about the same meaning, with nearly the same implications as (10) and (11): the addressee is doing something wrong with his political career and he is foolish enough to do so. We may argue that the reason for these coincidences is that the two expressions call for the activation of the same ICM, which we can call the Self-Annihilation ICM. This is a specification of the more general Annihilation ICM. Below we provide partial specifications of both ICMs:

The Annihilation ICM:

- (a) A person can kill another person (i.e. he may have control over the end of another person's existence).

- (b) Killing a person may be either intentional or unintentional.
- (c) On moral grounds the intentional killing of a person, whether the person is good or bad, is wrong.
- (d) The intentional killing of a good person is morally worse than the killing of a bad person.
- (e) Killing another person on purpose can only be justified when it occurs in self-defence or as the application of justice.
- (f) Killing another person unintentionally is foolish.

The Self-Annihilation ICM:

- (a) One is able to put an end to the existence of oneself (i.e. one has control over the end of one's own existence).
- (b) Killing oneself may be either intentional or unintentional.
- (c) Killing oneself, whether intentionally or not, is foolish.
- (d) On moral grounds the intentional killing of oneself, whether one is good or bad, is wrong.
- (e) Killing oneself on purpose can hardly ever be justified (i.e. there is no clear case of self-defence or of the application of justice).

The difference between (11) and (12) lies in the way the Self-Annihilation ICM is called up. In (12) this is done directly, but in (11) we have the mediation of the Grave-digging ICM which interacts with the DIVIDED PERSON metaphor, thereby providing additional contextual effects, as we shall presently see. This analysis will show that it is not accurate to postulate the existence of inconsistencies between source and target which may only be explained as part of the blend.

Let us start with a description of the Grave-digging ICM, which contains at least the following information:

The Grave-digging ICM

- A person digs a grave as an arrangement for another person's (expected) death. For example, if someone is going to be executed or is in the throes of death it might not be implausible to see someone else digging his grave. When the grave is finished it will be ready to receive the corpse.
- In that context, the digger will only finish the digging if the dying is sure to occur. Were the execution to be stopped or if the dying person were to be miraculously revived, the digging activity would probably be called to a halt.

In the metaphor the digger and the dying person from the source become different aspects of the same person in the target. This happens through the application of the Divided Person metaphor: the digger maps onto the subject, the rational part of the person that is unfortunately acting rather foolishly, and the dying person maps onto the self. Dying is the result of the person's misled activities (that is, of the mistakes made by the self), and digging corresponds to the foolish activity of the subject making arrangements to celebrate a success that will most likely never come. Also, the fact that the digging is still in process suggests that ultimate failure might be avoided

if the subject realized that what the self is doing is wrong and decided to stop the self. It is this aspect of the metaphor that explains why it is often used as a warning. Thus, the digging will only be finished if the misled behaviour eventually leads to downright failure. Compare:

(13) John has finally dug his own political grave

Although this sentence does not state that John has definitely ruined his political career, it strongly implicates that he has: if he has dug the grave and nobody has stopped him, it means that there was no need to stop the digging; therefore he has come to a point of no return (i.e. he has failed).

So, in this metaphor, death is not the direct consequence of digging, but a highly likely indirect consequence. The digger is dying because of other misfortunes, not because of his digging. His problem is that he does not realize how severe his situation is, and keeps on digging as if the grave were not for him. This formulation follows the Grave-digging ICM closely and is also compatible with the Extended Invariance Principle since it preserves the overall event shape (e.g. the causal, intentional, and frame structure) of the ICMs involved. If the interpretation we are giving is correct, the role of the blend is to become the repository where relevant portions of different ICMs can be put together and interpreted in combination. Any implication thereby derived will have to accord with the Extended Invariance Principle.

To sum up, in order to interpret metaphor (10) (which works under the correspondence GETTING INTO TROUBLE IS KILLING ONESELF), we need the combined projection into the blend of additional information from at least three ICMs: Grave-digging and Self-Annihilation, for the source; the Divided Person, for the target. Grave-digging in the source does not cause death directly but is a harbinger of death. This means that there is a grave-digger and a dying person. Since the target has only one person, this would render the mapping impossible. In order to avoid the incongruity, the person in the target is seen in terms of the Divided Person metaphor, whose activation is prompted by the wording of the metaphorical expression. Now, we have the foolish situation in which the subject keeps digging instead of trying to regain control over the self to avoid further harm. Since the harm might still be avoided, doing nothing about it is like seeking one's own death on purpose. This brings the Self-Annihilation ICM into the picture, with the information that killing oneself, whether on purpose or not, is foolish and hardly justifiable. As a result, the idea of foolishness associated with the person's overall behaviour is reinforced, thus differentiating (10) from other related metaphors, like (12) above, which also exploit the correspondence between making mistakes and killing oneself. So, the person in (10) is considered to be foolish not only on account of his misled activities (i.e. killing himself), but also on account of expecting success when imminent failure is evident and of making preparations for it (i.e. digging the grave while expecting somebody else's death rather than his own).

If the account given here is correct there are no mismatches between the source and target of metaphor (10). The causal structure is not inverted: the connection between grave-digging and death is simply established by implication. The inten-

tional structure is never violated: the digger in the source knows that he is digging a grave; similarly, the person in the target also knows (or at least thinks that he knows) what type of activity he is engaged in; however, he is unaware of the disastrous consequences of his actions. What the metaphor maps here is the relation existing between agent and actions in the source to the one existing in the target. Then, with respect to frame structure, it is not wholly accurate to say, as the authors do, that agent and patient are fused; we have seen that the agent is different from the patient in terms of the Divided Person metaphor (the subject does the figurative digging in order to have the self buried); they are only fused denotatively, but conceptually they are seen as different entities. Nor is there any necessary temporal connection between finishing the digging and dying but between finishing the digging and the imminence of death. This needs to be so in order to preserve the consistency between source and target in accordance with the Extended Invariance Principle. Note, in this connection, that in the source dying can precede the completion of the grave, but not in the target because of the impossibility of doing anything after death. Finally, neither in the target nor in the source is there any direct correlation between the depth of somebody's grave and his chances of dying. There is a different correlation: the deeper the grave the closer it is to completion and to being ready for its purpose (i.e., to receive the dead body of the person who has been acting foolishly).

Finally, the reason why there appear to be some inconsistencies between source and target is that our interpretation of the metaphor is procedural and automatic to such an extent that it becomes extremely difficult to analyse the way different ICMs are called to interact.

6. Conclusion

We have been able to show that a theory of blended spaces is compatible with the general requirements of a theory of production systems. We have seen that the systems may well provide the basis for inferential work, and that all conceptual activity associated with them is sensitive to general principles of conceptual consistency. This has led us to a vision of blended spaces as a by-product of the activity of working memory where matched productions retrieved from production memory are executed to yield pre-established combinations of ICMs. Since sometimes ICMs are called for to contribute to reasoning processes, this vision would provide a rationale for the study of contextual effects and discourse coherence. Finally, it has been shown that the idea that blended spaces contain mismatches and irregularities is largely illusory. Such phenomena are best explained as contextual implications triggered off by the converge and combination of the ICMs involved in the productions.

References

- Anderson, John R., 1983. *The architecture of cognition*. Cambridge, MA: Harvard University Press.
Anderson, John R., 1987. Skill acquisition: Compilation of weak-method problem solutions. *Psychological Review* 94: 192–210.

- Blakemore, Diane, 1992. *Understanding utterances*. Oxford: Blackwell.
- Fauconnier, Gilles and Mark Turner, 1994. Conceptual projection and middle spaces. La Jolla, Department of Cognitive Science, University of California, San Diego. [Internet document available from <http://cogsci.uscd.edu> and from <http://www.wam.umd.edu/~mturn>].
- Fauconnier, Gilles and Mark Turner, 1996. Blending as a central process of grammar. In: A. Goldberg, ed., *Conceptual structure, discourse and language*, 113–130. Stanford, CA: CSLI Publications.
- Lakoff, George and Mark Johnson, 1980. *Metaphors we live by*. Chicago, IL: University of Chicago Press.
- Lakoff, George and Mark Turner, 1989. *More than cool reason. A field guide to poetic metaphor*. Chicago, IL: University of Chicago Press.
- Lakoff, George, 1987. *Women, fire, and dangerous things: What categories reveal about the mind*. Chicago, IL: University of Chicago Press.
- Lakoff, George, 1990. The invariance hypothesis: Is abstract reason based on image-schemas? *Cognitive Linguistics* 1(1): 39–74.
- Lakoff, George, 1993. The contemporary theory of metaphor. In: A. Ortony, ed., *Metaphor and Thought*, second edition, 205–251. Cambridge: Cambridge University Press.
- Lakoff, George, 1996. The internal structure of the self. In: G. Fauconnier and E. Sweetser, eds., *Spaces, worlds and grammar*. Chicago, IL: University of Chicago Press.
- Newell, Allen, 1989. *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Ruiz de Mendoza Ibáñez, Francisco J., 1996. Blended spaces and the pragmatic approach to cognition. In: B. Penas, ed., *The intertextual dimension of discourse*, 233–244. Universidad de Zaragoza.
- Ruiz de Mendoza Ibáñez, Francisco J., 1997. Perspectives on metaphor. In: J.L. Otal, I. Fortanet and V. Codina, eds., *Estudios de lingüística aplicada*, 689–697. Castellón: Universidad Jaime I.
- Sperber, Dan and Deirdre Wilson, 1986/1995. *Relevance. Communication and cognition*. Oxford: Blackwell.
- Taylor, John R., 1989/1995. *Linguistic categorization. Prototypes in linguistic theory*. Oxford: Clarendon.
- Turner, Mark and Gilles Fauconnier, 1995a. Conceptual integration and formal expression. *Metaphor and Symbolic Activity* 10: 183–204.
- Turner, Mark and Gilles Fauconnier, 1995b. Blending and metaphor. In: Y. Shen and A. Kasher, eds., *Cognitive aspects of metaphor*. London: Routledge.