

Connectives are overt manifestations of a wide array of causality relations constructed when we use language. But we also build up elaborate causal chains in the absence of explicit markings. I analyze here cases of grammatically unmarked causal compression and decompression and suggest that their apparent lack of processing difficulty in spite of their great logical complexity deserves the attention of linguists and cognitive scientists.

1. Causal chains

Consider first the following, heard on the radio during a tennis match in which Martina Navratilova was playing.

(1) Martina is three points away from the airport

The meaning was clear: if Martina lost three more points, she would lose the match, be out of the tournament, and go home. Readers can verify for themselves that the very same sentence in a different context could have very different meanings. But for the tennis fan listening to the ongoing description of the match, there is no hesitation, and apparently no cognitive effort in selecting the appropriate interpretation. And yet that interpretation requires the construction of an elaborate causal chain which nothing in the grammar and words of the sentence explicitly signals.

Note first an obvious feature: there is no "literal" meaning for this sentence which could be processed first and then submitted to pragmatic adjustments. *Three points* is not a literal distance that could indicate where someone (Martina) is standing. The listener must engage directly in building connections that are compatible with the language form and with the context, although strictly speaking not imposed by either. This doesn't come as a surprise for cognitive linguists, who repeatedly stress that language prompts for meaning construction rather than symbolizing meaning directly.

What then is the meaning construction effected in the case at hand? Consider sentences related to the original:

(2) Martina is three feet away from the precipice.

(3) Martina is three points away from the end of the match.

(2) is a spatial description. (3) is a description of events, using the event structure metaphorical blend:¹ on a possible ordered sequence of events from the present to the end of the match, "going through" 3 losses of a point would cause the end of the match. Just as there are other paths in (2) than the three foot one to the precipice, there are other possible paths (sequences of events) in the tennis match than Martina's losing 3 points.

Of course, for any of this to make sense, some knowledge of tennis is required – the "tennis idealized cognitive model" with structure in terms of points, games, sets, matches, tournaments.² A sports event in which points are scored in succession by competitors can

be broken down naturally into events of scoring a point, and this in turn provides a metaphorical measurement unit (scoring a point) for the overall event. This particular measurement unit is different from time: one point can last much longer than another.³ Metonymic compression allows the word *point* to refer to the event of scoring a point, i.e. to the measurement unit. An additional twist observable in sentence (1) is that the *three points* would be points lost by Martina, not necessarily the total number of points leading to the end of the match.⁴ The subject of sentence (3), *Martina*, is positioned on the metaphorical path of events according to the event she is presently engaged in. If the end of the match is caused by Martina's losing, then sentence (3) is equivalent to (4):

(4) Martina is three points away from losing the match.

The structure of the tennis ICM allows the following causal chain to be constructed without additional explicit language:

lose three points --> lose set --> lose match --> be eliminated from tournament

So, (4) is equivalent to (5):

(5) Martina is three points away from being eliminated from the tournament.

A sports tournament ICM, that applies to tennis, provides the two alternatives:

win --> continue competition --> overall win --> champion --> award, congratulations, adulation ...

lose/be eliminated --> go home

As in any other ICM, the inferences obtained are not binding: an eliminated competitor can perfectly well stay on to watch the rest of the tournament, or travel on to the next tournament instead of going home, etc. The scenario in the ICM is a simplified stereotype.

Given the sports tournament ICM, (5) is equivalent to (6):

(6) Martina is three points away from going home.

One stereotypical scenario for "going home" is <going to the airport, getting on a plane, flying to hometown, getting off the plane, riding to one's house>. As explained in Lakoff (1987), the initial part of the trip (being at the airport) can stand metonymically for the entire trip:

Martina is at the airport --> Martina is going home

This provides the last causal link, connecting (6) to the original example (1). It allows one more twist: although the causal chain connects events that can have a certain "distance" from each other (points lost in the game), the ultimate point in that chain is

indicated by a spatial location (the airport), in harmony with the spatial preposition *away from*.

Although, I gave sentences (2) to (6) as support for the intermediate links in the causal chain, I don't mean to suggest that listeners need to go through the reconstruction of these linguistic forms in order to understand (1).

Rather, given the lack of any apparent cognitive effort to process (1) in context, it seems plausible that the relevant ICM's are highly activated and that as a result, the causal chains are highly accessible. The form that triggers the required meaning construction, namely sentence (1), yields a superb compression in context: from the single noun phrase *the airport*, the entire causal chain of events is "automatically" decompressed thanks to the activated ICM's:

**lose three points --> lose set --> lose match --> be eliminated from tournament
--> go home --> <pack suitcase --> go to airport --> be at airport --> get on plane
--> fly to hometown --> go to house>**

To sum up, then, the "points" stand metonymically for the events of losing a point. Three such events would lead to defeat. The events are on a metaphorical spatial scale to which the tennis player gets mapped. On that scale, the player is metaphorically at a spatial distance of "three points" from the end of the match which would mean defeat. A metonymic chain takes us from the end of the match to defeat, then to exclusion from the rest of the tournament, then to returning home. The airport (a place) stands metonymically for an event (flying home) that starts in that place. Through the chaining, "flying home" links to "leaving the tournament" which links in turn to "losing the match," itself caused by "the three lost points."

Strikingly, very little of this is indicated by the linguistic structure itself. It is constructed by means of the cognitive models that we have for games, tennis, tournaments, travel and by applying to them the appropriate mappings. Importantly the successive mental leaps are not coded in the output signal and yet they are fully reconstructed by the listener, apparently without any unusual cognitive effort, as long as the relevant ICM's are available (context effect: the tennis situation has maximum salience). Asked to come up with a meaning for this sentence in isolation, subjects are able to imagine a variety of scenarios, but unsurprisingly do not come up with the original announcer's meaning relating to tennis.

Finally, compare the following:

(1) *Martina is 3 points away from the airport.*

(7) *Svetlana is an inch away from the podium.*

[high jump: a final jump of one inch more than her previous ones will put Svetlana in the top three]

(8) *The podium would mean a lot to Svetlana.*

(9) **The airport would be a disaster for Martina.*

The type illustrated by (7) is a variant of the type illustrated by (1). But there are some noteworthy differences. In (7) and (8), *the podium* can function as a metonymic target

for the event of *winning*. The anomaly of (9) shows that this is not the case for *the airport*. Even though, as we showed above, there is a long reconstructed causal chain that includes *losing* and *the airport*, that is not enough to sanction a full metonymy: (1) is indeed equivalent to *Martina is 3 points away from losing*, but (9) cannot be used to mean that *losing would be a disaster for Marina*. A possible reason could be that *podium* is a component of the competition frame directly associated with *winning*, whereas *airport* is part of a completely different *travel* frame.

The Martina example is just a particularly visible case of massive causal compression in context. It is not amenable in standard ways to psychological experiments based on reaction time or ERP's because the complexity of the meaning construction is not mirrored by a corresponding processing difficulty or syntactic complexity. I have given a plausible explanation of what is going on, but no account of the actual psychological/neurobiological story for how the brain is able to accomplish all this. It is a research program that deserves to be pursued.

2. Causal compressions in conceptual blends

The compression of vital relations in conceptual blends has been studied in some detail.⁵ Cause-Effect is routinely compressed and decompressed. Again, no explicit causal connectives need appear, but Cause-Effect, as a "vital relation" is readily reconstructed.

In the integration networks, Cause-Effect is often compressed into Uniqueness. For example in the celebrated Bypass example,⁶ the children's poor education will cause their low competence as surgeons years later. In the blended space, the children are merged with the surgeons they might become, and their ignorance becomes the surgeons' incompetence. In Eve Sweetser's analysis of rituals – the baby in the stairs, the buffalo hunt – a causal relation between the satisfactory accomplishment of the ritual actions and success for the corresponding "real" events (the baby's life, the hunt) is compressed: the ascent up the stairs "is" the person's life, the ritual dance "is" the hunt.⁷

Compression of long diffuse causal sequences to human scale single intentional action is equally striking, as in the example where the veto of a foreign aid bill by a politician is described as:

(10) *He's snatching the food out of the mouths of starving children.*

(11) *He's killing innocent children.*

The literal falsity of the emergent descriptions in the blended space (*snatching food*, etc.) is in no way an obstacle to the compression. In fact, it seems to make it more powerful.

Here is how such cases work. One input has the actual vetoing of the foreign aid bill by the politician. The other input has the human scale event of a grown man snatching food from a starving child. In the "veto" input, pattern completion of a standard scenario operates to add an extensive causal chain of the form: **veto --> no foreign aid --> no money to buy and transport food --> no food transported to needy children --> the**

children have no food --> the children starve. A (third) counterfactual input is constructed in parallel: **bill passed --> foreign aid --> money to buy&transport food --> food transported to needy children --> children eat the food --> children survive.** All three inputs are selectively projected to the blended space: the framing is the human scale snatching of food. The food snatched is projected from the counterfactual space (without the veto, the child would have this food), the elaborate causal chain in the veto input, with the politician as initial causer and the child without food as ultimate consequence is compressed to the single event causal chain of snatching the food from the child (politician causes the "counterfactual" food to be "taken away" from the child). Notice that the human scale input of literally snatching food also has a covert counterfactual of its own: without the snatcher (cause), the child would keep the food and eat it.

To be able to decompress (10), the listener must have contextual access to the veto input, be able to extend it by means of an elaborate causal chain, build the relevant counterfactual, and construct the integration network as outlined above. As in the Martina example, but this time through a different process, it is striking that this elaborate causal decompression is automatic, produces no cognitive perplexity, and apparently no cognitive effort. The complexity of the meaning construction cannot be measured by assessing processing difficulty. Novel experimental designs that could address this remarkable human capacity would be welcome.

Here is a remarkable piece of data discovered by Eran Agmon, which uses the very same principles in more spectacular fashion. This is a poster used during World War II to encourage Americans to save gas (needed for the war effort) by car-pooling.



The structure of this conceptual blend is fairly straightforward. One input is the familiar human scale scenario of helping someone by giving them a ride. The other input has the elaborate causal chain **riding alone --> use too much gas --> waste gas --> harm the community that needs gas --> harm the country at war --> help the enemy --> help Hitler**. As in the veto example, a third input is also covertly set up: **car-share --> save gas --> help country at war --> hurt the enemy --> hurt Hitler**.

The blended space is framed by the human scale scenario of the "giving a ride" input. Hitler is projected from the "riding alone" input into the role of person being helped. As in some other highly visible blends, Hitler takes on "ghost-like" features that help to prompt the decompression and block any literal interpretation. A covert counterfactual for the human scale input is perhaps transporting other people (car-sharing), countrymen that you help, and who occupy all the seats in the car, preventing Hitler's ghost from sharing the ride.

Clearly, a massive causal decompression must take place in order for all this to work: the causer (driver and single occupant) in the war input is mapped onto the driver with an invisible passenger in the human scale input. The elaborate causal chain of driving along to helping Hitler (in the war input) is mapped onto the single causal event causal of helping the passenger in the human scale input. In the blended space, the causal chain is completely compressed to the driver helping the invisible Hitler.

3. Instant build-up of multiple spaces

To show in some detail our capacity to construct elaborate causal structures in order to understand what looks like a simple language form (and to highlight the necessity of doing so), I will take another look at Coulson and Pascual's example (in Coulson and Pascual (2006)):

Would I kill my daughter so I could walk again?

The speaker is Steve Johnson and the issue is the appropriateness of stem cell research:

"Among the parents was Steve Johnson, a paraplegic who, with his wife Kate, adopted an embryo whom they named Zara — now a little girl in a pink, flowered dress and blond curls playing near her father's wheelchair. Johnson described the years of pain, high medical costs, and limited mobility he'd endured after a bike accident 12 years before. 'My soul aches for a cure for my paralysis,' he said — but not at the cost of a child's life. 'Would I kill my daughter so I could walk again? Of course not. Then why do we think it is okay to kill someone else's kid?' he asked."

In order to understand Steve's (rhetorical) question, an elaborate integration network needs to be constructed, with two alternative paths (one counterfactual) and multiple causal links (not marked by the grammar of the example sentence).

One path (the actual situation) corresponds to the adoption of the embryo. A network is constructed in which Steve and the embryo are first mapped onto a future situation of Steve with child Zara, that fits the father-child frame (F, C) [fig. 1].

For ease of reference, I will call the spaces <e>, <S e>, <S Z>, and <F C>. The first three are linked chronologically. <e> has the embryo by itself (perhaps frozen). <S e> has Steve being given the embryo.⁸ <S Z> is the imagined future space in which the

embryo will become a little girl, Zara. $\langle F\ C \rangle$ is the space with the partial kinship frame, father child. **e** and **Z** are linked by the identity connector.





Fig. 1

Spaces $\langle S\ e \rangle$ and $\langle S\ Z \rangle$ match and can be blended, yielding the network in fig. 2:

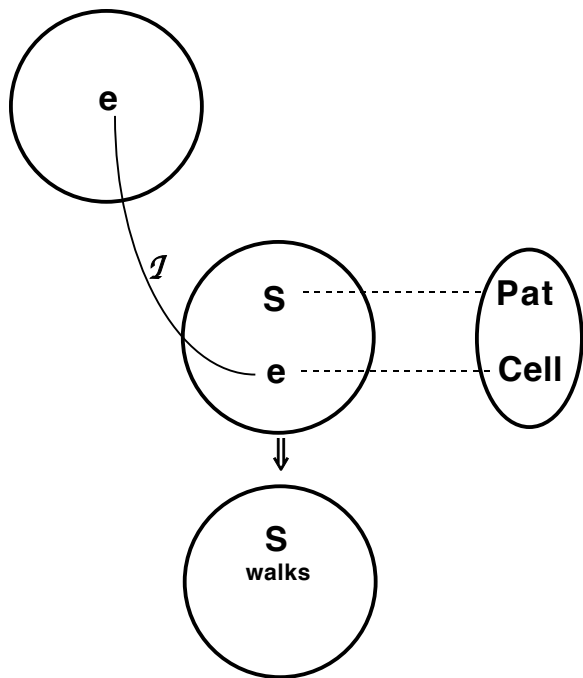


Actual Scenario

Fig. 2

In the blended space, the embryo is the child and Steve is the father. Identity (of embryo and Zara) have been compressed into uniqueness, and so, naturally, the embryo has the name Zara, it is a child, and it can be adopted.

In the second alternative path (counterfactual scenario), Steve is mapped onto the patient (in the medical frame Patient/Cells) and the embryo is mapped onto the role "cells" in that frame. Causality here is between the patient frame (in which the cells cure the disease) and the next mental space (in time) where Steve can walk (having been cured of his paralysis). See Figure 3:



Counterfactual Scenario

Figure 3

In the resulting mental space, where Steve walks, there is no counterpart for **e**, the embryo. That counterfactual space can be blended with the blended space of the actual scenario, to produce the conceptualization of a "missing" Zara:




Fig. 4

The killer/victim frame is recruited by pattern completion to produce the final causal sequence (fig. 5):

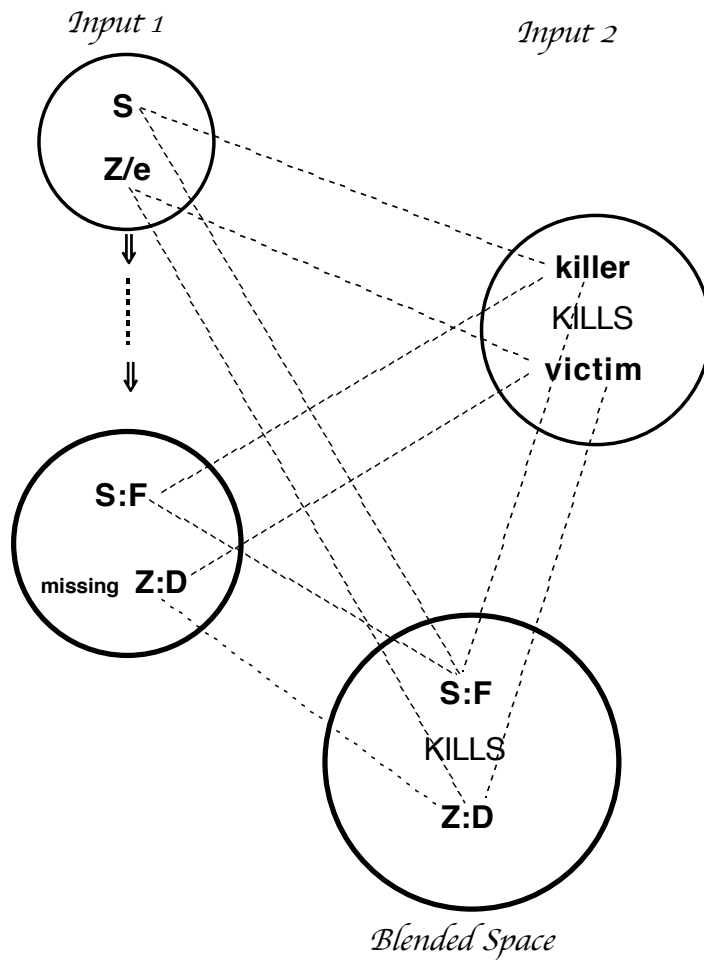


Fig. 5

In the final blended space, Steve kills Zara. In other words, the complex causal sequence (in the counterfactual) that leads to missing Zara is compressed into the single intentional "human scale" direct action of killing. This is similar to the politician's veto, except that the "victim" is constructed through far more elaborate blending. Steve, as the initial causer in the complex causal sequence that leads to Zara's "disappearance," is projected to the causer (killer) in the compressed human scale blend. Importantly, the killing frame cannot be imposed directly on the patient/cell frame, since the cells are biologically alive. But of course, through identity chaining of **e** in the patient/cell frame to Zara in the actual scenario, "killing" is projected back to **e** in the alternative counterfactual scenario.

Through another familiar (specific for generic) compression, Zara can then stand for any embryo; all embryos are potential daughters, hence actual daughters of potential fathers, hence of actual fathers, and any stem cell treatment is a killing of someone's daughter.

An apparently simpler analysis would just say that in Steve's belief system, embryos are children, and using them for any purpose other than the normal evolution from embryo to fetus to child coming into the world is therefore tantamount to killing. But this is really the same analysis without the intermediate causal steps and compressions fully spelled out. The "actual" path, with the adopted embryo, Zara, and the visible child Zara is enabled by the powerful conceptual blend of Fig. 2. The causal link between Zara in the actual scenario, and her disappearance in the counterfactual one is what allows the projection of the "killing" frame.

Metaphorical uses of "killing" use a similar build-up of actual vs. counterfactual in very standard, non controversial ways. To say that "you killed my project," is to set up two alternatives, one in which you act, and the project "aborts," and one in which you don't act and the project develops into a concrete result. Your action, intentional or not, can then fit "killing" in the killing frame. Compression allows blending of a "project" with its result. *My current project is a book on dinosaurs* is an appropriate statement before any writing starts, or at any stage during the writing, or when the completed manuscript is submitted for publication.

4. Overall goals and issues

The overall problem raised by causal compression phenomena is the following: elaborate space configurations are set up and processed in the absence of explicit grammatical markings for the frames, spaces, and connections needed. Language guides us in certain directions and constrains the configurations, but it does not directly specify the causal chains that need to be constructed in order for meaning to emerge. This is counter-intuitive. It is widely assumed in classical semantics that overt connectives are responsible for the logical structure of discourse. The deeper mechanisms that create logical structure from non grammatical cues deserve to be empirically investigated and theoretically accounted for.

¹ Lakoff and Johnson (1999), Fauconnier and Turner (in press).

² Listeners can have very different icm's/frames, accurate or not, detailed or sketchy. With no knowledge of tennis at all, a listener using only the competition icm and guessing that tennis players score points, could construct a meaning poorer than the intended one but still compatible with it.

³ The scoring system in tennis has the peculiarity that points don't just pile up. At deuce (40 all) a new system kicks in, and there are countless additional substructures, that the listener may know or be activating to different degrees.

⁴ So for example, if the score is 6-4, 5-2, 15-love, in favor of Martina's opponent, sentence (1) is appropriate, but there are several ways in which Martina could lose 3 points and the match. Three points lost in a row would do it, but also for the last game of the match: 15-love, 15 all, 30-15, 30 all, 40-30, game. In that case, 5 points are played, and Martina loses three of them.

⁵ Fauconnier and Turner (2003), Fauconnier (2005).

⁶ This example is discussed in the references of footnote 5.

It is an ad meant to persuade readers to help in the fight to raise standards in American schools. It shows three doctors in an operating room, who seem to be looking in the direction of whoever happens to be reading the ad. The headline is a voice introducing the doctors to the reader, who is also the patient. It says, "Joey, Katie, and Todd will be performing your bypass operation." The only odd thing about this scene is that Joey, Katie, and Todd are about seven years old. The body of the ad explains that doing anything sophisticated, like practicing medicine, requires sophisticated learning, but that America's kids are getting dumbed-down curricula. They won't understand chemistry, or laser refraction, or immunology, so they won't be good doctors, and the public, personified by the reader, will be at risk. Specifically, Joey, Katie, and Todd will operate on you, and you will probably die. Therefore, you should help get standards raised.

⁷ Sweetser (2000).

⁸ The space diagrams reflect what is focused on in the meaning construction: Steve as father and potential patient. Of course, a lot more is going on in this story, in particular Kate's role as adoptive mother and birth mother.

References

Coulson, Seana. 2001. *Semantic Leaps: Frame-shifting and Conceptual Blending in Meaning Construction*. New York and Cambridge: Cambridge University Press.

Coulson, Seana and Esther Pascual. 2006. "For the sake of argument. Mourning the unborn and reviving the dead through conceptual blending." *Annual Review of Cognitive Linguistics*, 4, 153-181.

Fauconnier, Gilles. 2005. "[Compression and Emergent Structure.](#)" In S. Huang, ed. *Language and Linguistics*. 6.4:523-538

Fauconnier, Gilles & Mark Turner. 2003. *The Way We Think*. New York: Basic Books.

Fauconnier, Gilles and Mark Turner. to appear. ["Rethinking Metaphor."](#) In Ray Gibbs, ed. *Cambridge Handbook of Metaphor and Thought*. Cambridge University Press.

Lakoff, George. 1987. *Women, Fire, and Dangerous Things* (Chicago: University of Chicago Press), case study 1, pages 380-415.

Lakoff, G. and M. Johnson. 1999. *Philosophy in the Flesh*. New York: Basic Books.

Sweetser, E. 2000. Blended spaces and performativity. *Cognitive Linguistics* 11 (3/4): 305-333.