The lazy mindreader : a humanities perspective on mindreading and multiple-order intentionality

Duijn, M.J. van

Chapter 4

Viewpoint Packages: linguistic tools for communicating and processing complex "thoughtscapes"

A famous athlete shoots dead his girlfriend at night in his house. The next morning, readers of newspapers all over the world find out that a spokesperson *states* that police officials *declare* that the athlete *claims* that he *thought* that he was shooting at a burglar, while the responsible police detective *claims* that he *knew* it was his girlfriend. Clearly, this complex "thoughtscape" is not represented in the news media in this way, using multiply-embedded sentences. I demonstrate that the representation of the involved mindstates relies substantially on lexical units implying viewpoint layers, such as *allegedly*, *accidentally*, *mistake*, and *to confirm*. I introduce and discuss the concept "viewpoint package", building on an existing framework that deals with meaning construction more generally (Dancygier's narrative spaces) and on one relevant account focussing on a particular part of speech (Vandelanotte's framing adjectives). Viewpoint packages allow for efficient coordination of multiple interrelated viewpoints in discourse, while regulating (audience's perception of) the commitment various discourse participants make to parts of the presented contents. I end by tentatively suggesting that viewpoint packages qualify as "tools for thinking": knowing their use in language may serve not only communication, but also support cognitive processing of complex thoughtscapes.

4.1 Introduction

On Valentine's Day 2013 a dramatic event made headlines all over the world: early that morning, a famous Olympic athlete shot dead his girlfriend through the bathroom door. The question that immediately perturbed everyone was whether it was murder, or a fatal accident: the athlete claimed that he had mistaken her for a burglar, while the police arrested him on the charge of having killed her wittingly and on purpose.[56] The crucial point of debate in this case was thus not whether he shot her or not, but whether he knew he would be doing so when he pulled the trigger.

This means that the choice between accident or murder coincides completely with the construal of the athlete's mindstate at a particular moment during the night of the shooting. Ultimately, this construal was made by a judge in court, based on information from forensic research, interrogations, witness reports, and so on. But besides that, and from the very first day after the incident, thousands of people have made such construals for themselves, mostly relying on cues presented by the news media. Although the main focus is clearly on what Pistorius thought, knew, and intended during the night of the shooting, most of these cues do not directly concern the athlete's mindstates. Rather, they add up to a complex "thoughtscape", a network of mutually embedded and interlinked viewpoints that are in some way relevant to the case: news media *suggest* that various sources *report* that the athlete *claimed* that he *thought* that his girlfriend was still in bed and not behind the bathroom door when he fired his gun. At the same time they *report* that a spokesperson *declared* that police officials *considered* it to be likely (or at least to a high degree possible) that the athlete *did know* that his girlfriend was behind the bathroom door, and that he thus *intended* to kill her. The news media also *report* what witnesses *claim* to have heard or seen, or what various sources *report* that family members *have declared*.

[56] Background to the case: South-African athlete Oscar Pistorius shot and killed his girlfriend Reeva Steenkamp on February 14th, 2013. Pistorius is a sports icon also known as "the Blade Runner"; his legs were amputated and yet he became a sprinter using carbon-blade prosthetic legs. In the aftermath of the killing, news media have frequently reported details of the court case, police investigations, personal life of Pistorius and Steenkamp, etcetera.

This entire thoughtscape is somehow represented in the headlines and articles about the case. Partly this is done using means of viewpoint coordination that are well-accounted for in the literature on speech and thought representation, such as indirect speech or free indirect speech. Consider, for example, the following passage from a press release on the morning after the shooting:

(1) Athlete Oscar Pistorius allegedly accidentally shot dead his girlfriend at his house in Pretoria on Thursday morning, Beeld.com reported.
(SAPA, 'Oscar Pistorius shoots girlfriend: report', 14 February 2013)[57]

Using a form of indirect speech, the propositional content of the reported clause "Athlete Oscar Pistorius…Thursday morning" is attributed to the perspective of *Beeld.com*. Yet there are more viewpoints coordinated in this sentence: through the adverbs "allegedly" and "accidentally", readers of (1) also learn that a source other than *Beeld.com* claimed that Pistorius killed his girlfriend without intending to. The information that Pistorius shot his girlfriend is clearly the critical bit, but it comes embedded in a complex of viewpoints that could be depicted as follows:

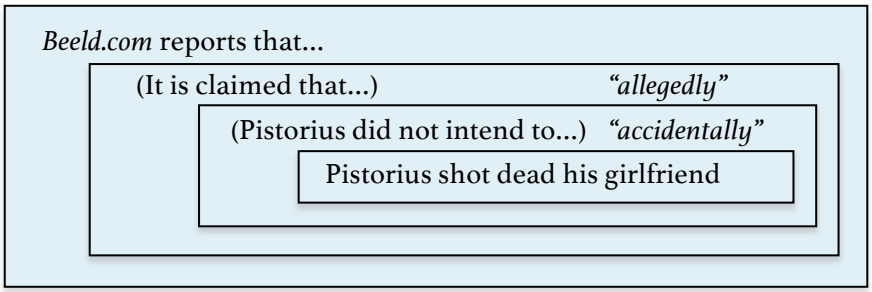

Figure 1

[57] The reports published on 14 February 2013 were retrieved on 21 February 2013 using Lexis Nexis (http://academic.lexisnexis.nl).

In this chapter I will contrast the two ways in which viewpoint layers are introduced into discourse as exemplified here: through a form of indirect discourse on the one hand, and through the use of single words (or lexical units) implying one or more viewpoint layers, such as *alleged(ly)*, *accident(ally)*, *mistake(n)*, *to confirm*, on the other hand. Various aspects of the working of such words have been studied in Vandelanotte's research on "framing adjectives" (2002; 2007) and in approaches dealing with meaning construction more generally (framing, blending, mental space theory).[58] Fauconnier (1997) makes the case that newspaper articles most of the time rely heavily on background knowledge pre-existent in readers. Given that this knowledge is for a large part structured in frames, journalists can activate a whole network of relevant information in readers by mentioning only a few well-chosen cues. This information is then available for interpreting the news details given in the article. The processes of meaning construction focussed on in this chapter could be considered a viewpoint-specific subcategory of framing in this sense. I propose to distinguish a set of words, across part-of-speech boundaries, capable of what I will refer to as "holistic" viewpoint coordination. The label that I propose for this set is *viewpoint packages*. Viewpoint packages, such as *alleged(ly)* and *accident(ally)*, activate readers' background knowledge related to viewpoint, which can then be assimilated with contextual details.

In addition, in my final section I will address the relation between linguistic discourse and cognitive processing by suggesting that viewpoint packages can be seen as instantiations of what Dennett calls "tools for thinking" (2000; 2013): knowing their use in language may alleviate the burden of handling multiple perspectives cognitively.

[58] It should be noted that "framing" in "framing adjective" goes back to a different tradition of using this term: see McGregor (1997: esp. 66-67).

4.2 Viewpoint, mindstates, and the human ability to read minds

"Viewpoint" forms an object of study situated on a busy crossroads of disciplines. Whereas concepts such as "mindstate" and "intentionality" go back a long way in philosophy of mind, "viewpoint" and "perspective" have been traditional topics in the study of language and literature. During the past decades, psychologists and cognitive scientists have taken up the subject, and in recent years there has been an increasing number of attempts to bring insights from different angles together. The approach proposed in this chapter can be seen as an attempt to pursue this line.

In what follows, by a "mindstate" or "intentional state" I mean a particular belief, intention, desire (etc.) held by an animate individual with respect to a state of affairs in the outside world (much in the fashion of Dennett's *intentionality*; 1987). I use "viewpoint" and (in this chapter synonymously) "perspective" as indicating the broader, overall take a person, group, or institution (e.g. a newspaper) has on a certain part of that world. Intentional states or mindstates are thus, as it were, atomic "snapshots" of a subject's relation to an object; a perspective or viewpoint comprises the broader total of an actor's subjectivity of which intentional states are isolated parts.

By definition, neither mindstates nor viewpoints/perspectives held by others can be accessed directly, but they can be appraised through an inferential process based on behavioural cues (including linguistic utterances) and immediate circumstances. This latter process is often referred to as "theory of mind" or "mindreading" (for an overview see Apperly, 2011). Several cognitively-oriented linguists and literary scholars have pointed out that the process by which language users form an understanding of the viewpoints and mindstates of people or characters referred to in discourse, can be seen as a special case of theory of mind or mindreading.[59] In the physical presence of an interlocutor (say, when speaking with John), we have direct access to verbal and

[59] See e.g. Verhagen (2005). For fiction see Palmer (2004); Budelmann and Easterling (2010); Dancygier (2012); Sluiter, Corthals, Van Duijn, and Verheij (2013); Cefalu (2014); see also Chapter 2 and 3.

non-verbal cues that may guide our inferential process regarding John's mindstates. However, when John is not physically present but referred to by Mary in a conversation we have with her (or a story narrated by her), a similar inferential process can be triggered through cues in Mary's speech that enable us to construe John's mindstates or viewpoint. Applied to our object of study in this chapter, understanding a thoughtscape on the basis of news reports is thus seen as a form of mindreading in which textual cues guide the inferential processes within readers.

In the next section I will investigate this process more closely, while bringing elements from a variety of frameworks together. As a first step in my analysis I will distinguish between two ways of introducing viewpoint complexity into discourse: *compositionally* versus *holistically*.

4.3 Constructing complexity*

4.3.1 Compositional complexity: a literary example

Regarding relevant approaches to meaning construction, in particular mental-space theory, framing, and blending, we will mostly rely on Dancygier's narrative-centred synthesis offered in her 2012 book *The Language of Stories*. Her narrative-spaces framework forms a useful tool for analysing how texts can represent a complex thoughtscape. The central issue in this framework is how cues at the micro-level of a text ultimately yield an increasingly rich and complexly structured story at the macro-level. Dancygier builds on mental space theory and offers a related though newly devised core concept: the *narrative space* (see Dancygier, 2012: esp. chapter 2; for mental space theory see Fauconnier, 1997; 1985). Like a mental space, a narrative space is a hypothetical subdivision in a language user's mental activities, prompted by linguistic expressions, and used in the process of online meaning construction. Narrative

* Just a reminder of what is explained in the Reading Guide at the beginning of this thesis: the basics of the narrative-spaces framework are also discussed in Chapter 2 and 3, though with slightly different emphases. This is due to the fact that the chapters are written as separate papers, aiming at slightly different audiences.

spaces are each characterised by a particular set of features such as time, space, cultural norms, language spoken, or participants involved in either narration (narrators/focalisers), action (characters), or both (Dancygier, 2012: 35-37). These features are open to further elaboration by all kinds of local linguistic choices, such as usage of elements that coordinate viewpoints (for example complementation constructions or, as I argue, lexical units such as *allegedly* and *mistaken*), grammatical features such as tense or modality, usage of pronouns, typographical cues (such as quotation marks), etcetera.

Construction and elaboration of narrative spaces take mostly place through processes generally known as *framing* (Fillmore, 1985; for an overview see Cienki, 2007) and *blending* (Turner and Fauconnier, 1995; Coulson, 2001). The basic idea is that the background knowledge used by readers to interpret a text is structured in frames. Particular linguistic items used in the text activate a frame in its entirety, even if they relate to it only indirectly and in an unpredictable manner (cf. also the notion of "frame metonymy" as discussed in Fauconnier, 1997). Using this principle, writers of news articles can evoke rich meanings while providing only a few cues. For example, if a headline provides the information that a neighbour heard "non-stop shouting" coming from the Pistorius home during the hours before the shooting, nothing more needs to be said to evoke the frame of a *fight* between lovers, the sort of context in which one can imagine that things went *out of hand*, leading to a directed *attack* and hence a case of *murder* rather than a tragic accident (see also Section 4.3.6 below). In the current chapter I focus on a subcategory of this more general phenomenon, specific for viewpoint coordination and linked to a category of lexical items I refer to as viewpoint packages.

The general idea of blending is that two inputs with an established conceptual structure and content are integrated into an emergent *blended space* or *blend*. The blend has properties inherited from the inputs, as well as new structure and content of its own. Once a narrative space has been prompted and possibly further structured and enriched, it can as a whole be blended with another narrative space. The result of all the construction and blending processes is what Dancygier refers to as the *emerging story*: this is what a reader understands after having read and processed the text up until a particular point. The emerging story is thus a "moving end result".

Although Dancygier's framework is primarily designed to analyse narrative fiction, it fits our object of study in this chapter well. The ultimate question in the Pistorius case clearly is: was it an accident or murder? News reports provide a variety of micro-level cues that enable readers to put together a version (or rather, multiple competing versions) of the story covering the hours, days, or even weeks before the shooting took place. Clearly, the genesis of the story is different: instead of an author inventing it and "feeding" it gradually to the readers, the journalists themselves build their understanding of the case on various sources of facts and opinions. My focus here is on the linguistic choices these journalists make when subsequently feeding what they have understood to the readers of the newspapers. The narrative-spaces framework can be used to analyse how the journalists' linguistic choices at the micro-level of news reports relate to the formation of an increasingly rich and complex emerging story understood by the readers of the news.

One of the fictional examples Dancygier uses to introduce the concept of narrative spaces comes from Eggers' novel *A Heartbreaking Work of Staggering Genius* (2000). The main character Dave worries about having left his little brother Toph at home with a baby-sitter. While he is driving in his car, anxious thoughts cross his mind:

(2) I will come home and the door will be open, wide. The baby-sitter will be gone [...] Blood on the walls [...] a note to me [...] There will be a hearing, a trial, a show trial –

How did you come to meet this man, this baby-sitter?

We found a posting, on a bulletin board.

And how long did your interview of him take?

Ten, twenty minutes.

(based on Dancygier, 2012: 38, citing Eggers 2000: 126; italics in Eggers' original)

The main narrative space here is a period in Dave's life. Within this main narrative space, particular micro-level language phenomena prompt the construction and structuring of additional narrative spaces. Examples of such

phenomena are the choice of pronouns ("I", "you", "we"), choice of verb tense and modality ("will come", "will be", "did", "found"), and the use of italics ("*How did you...babysitter*"). The result is an emerging story that contains an increasingly rich and complex structure. Within the basic setting of Dave's life, the moment captured by the passage in (2) is a car ride to San Francisco. Within the setting of that car ride Dave imagines the scenario of coming home to where he left his little brother and finding blood on the walls, followed by an imagined trial in which he is being interrogated about the babysitter, realising (but this is not spelled out) that he did not find him through very reliable channels and did not take much effort to check on him either.

The emerging story at the end of (2) is indeed a blend of several narrative spaces. At the point of the interrogation, the story has parts of the structure of all these spaces: Dave is still in his car driving to San Francisco, and at the same time, through the layers of imagined events, there are his empty house with blood on the walls and the interrogation in court. However, readers will clearly not conclude that little brother Toph is "really" dead and Dave is "really" facing a trial. This is because they have sequentially read through the process in which the structure was built up one layer after the other, prompted by language phenomena that not only added content and structure, but also provided information about *how* this content and structure had to be integrated in the emerging story. The modal verbs at the beginning of (2) ("I *will* come home", "the door *will* be open", etc.) and other formal choices (such as the use of ellipses and enumerations, comma's, italics, etc.), along with the absence of any direct evidence of a crime (such as a phone call from the police that Toph was found dead), signal to the readers the status of the presented content: the blood, the note, the trial, etcetera, must be products of Dave's imagination. But what the readers *do* conclude from the fact that Dave is imagining doom scenarios, is something about Dave's overall mental condition: he probably has hysteric or paranoid tendencies, in particular when it comes to his little brother. However, overarching terms such as "hysteria" or "paranoia" are nowhere mentioned in the text. In other words: the constituents (imagined terrifying scenarios) are given, but the "whole" (hysteria) is omitted and left for the reader to construe. This is an instantiation of what I refer to as *compositional* complexity: the components that constitute the structure of the emerging story are all spelled

out in the text and each adds a single part of the total complexity, while the construction of an overall diagnosis of what is going on is left to the reader.

4.3.2 Compositional complexity: an example from the news

Similar to this, viewpoint layers are spelled out explicitly, hence complexity is construed compositionally, when a form of sentence embedding is used to coordinate perspectives. Consider again the sentence cited in the introduction, here repeated for convenience:

> (3) Athlete Oscar Pistorius **allegedly accidentally** shot dead his girlfriend at his house in Pretoria on Thursday morning, *Beeld.com* reported.
> (SAPA, 'Oscar Pistorius shoots girlfriend: report', 14 February 2013).

In the introduction I have stated provisionally that the first layer (*Beeld.com*) is related to the three others using "a form of indirect speech". However, in the literature there is no consensus on the question whether this should be considered a case of Indirect Speech (IS) or Free Indirect Speech (FIS). The expression in (3) as a whole is a specific syntactic pattern consisting of a reported clause followed by a reporting clause (underlined) on which the former is grammatically dependent. The clauses are not paratactically related, and the reporting clause is not grammatically complete without the reported clause (*"*Beeld.com* reported."). Consequently, the construction is one in which the reported situation is embedded in the point of view indicated in the reporting clause, and in this respect it is comparable to complementation constructions (although the different syntax also suggests semantic differences). Jeffries and McIntyre (2010: 89) would categorise it as IS, arguing that a sentence such as "The weather was nice, John said" only differs from "John said the weather was nice" in the order of the clauses, and the latter is clearly IS. Leech and Short (2007: 276) also mention "inversion", but at the same time have more of an eye for the differences when calling this pattern "Janus-like", "somewhere in between IS and FIS". Toolan (2006: 703) observes that the complementiser "that" cannot be inserted in preposed reported clauses and that they sometimes clearly require independent clause syntax ("Could he

accompany her home, he asked"; not *"He could accompany her home, he asked"). He therefore concludes that preposed reported clauses are more like prototypical FIS, and might even be considered FIS if it were not for the inquit-formula (for more discussion and examples see also Vandelanotte, 2009; 2012).

Given the presence of the inquit-formula (indicating a construal of the reported clause as in some way dependent), while at the same time agreeing with Toolan that this is clearly not a complementation construction, I propose to consider (3) as a relatively autonomous embedding construction, which I will refer to as *inquit-construction*. The important similarity between an inquit-construction and "classic" indirect speech using complementation, is that both coordinate two viewpoint layers that are spelled out in the text separately. Readers are cued to interpret one part of the sentence's content as being part of one layer (or narrative space), and another part as belonging to a different one. This is what makes it a form of compositional complexity: as in the Eggers-example, the constituting layers are spelled out in the text, whereas the overarching "whole" is not referred to explicitly. The grammar of English thus contains a family of constructions, including complementation and the inquit-construction, for cuing this kind of compositional complexity.

4.3.3 Holistic complexity: viewpoint packages

By contrast with compositional complexity, viewpoint packages introduce complexity holistically: the wholes are given in the text, while the underlying constituents remain for the reader to construe or unpack – if the context so requires. I argue that words such as *allegedly* and *accidentally* (bold in (3)) are holistic prompts to readers to imagine viewpoint layers that are not spelled out in the text. Vandelanotte discusses *alleged* as what he calls a "framing adjective", capable of "set[ting] something apart as belonging to a 'second-order' reality, viz. the reality of another's discourse" (2007: 360; referring also to McGregor, 1997: chapter 6). Other examples of framing adjectives are *so-called*, *pretended*, *purported*, or *supposed*. Vandelanotte suggests that "their shared reportative-evidential meaning can tentatively be glossed" as follows: "[framing adjective] X = stated by someone, but not the speaker, to be X" (2007: 368). Developing this line further, I suggest that Vandelanotte's "tentative gloss" reflects a *topology*, a

piece of conceptual structure evoked by the adjective. In the case of *allegedly* this conceptual structure looks in some respects similar to the one set up by a complementation construction or inquit-construction: one part of the related content must be attributed to the speaker and another to a third party. I will argue that such topologies implying coordination of viewpoints are not exclusive to framing adjectives; words that can function in this way can be found across several parts of speech. As said, the label that I propose for this type of words is *viewpoint packages*.

Looking at *accidentally* in this way, the following topology can be postulated: there is a particular outcome of an action, which differs from the one intended by an involved agent. In the actual discourse concerning the Pistorius case this topology is mapped onto situation-specific details, such as the athlete being in his house with his girlfriend at night and shooting her. When processing sentence (3) above, the reader will assimilate (through blending) the topology of *accidentally* with these details and take it that the athlete shot dead his girlfriend, but had not intended to do so.

In the case of *allegedly* it is given in the topology that a source different from the speaker asserts the information under the scope of this adverb. The details provided by the context in (3) only partly elaborate the topology and leave the identity of this other source open: it could be the police detectives or someone else who is in a position to make the claim that Pistorius (claims that he) shot dead his girlfriend accidentally. What content exactly is affected by *allegedly* depends on the interpretation of its scope: if only *accidentally* comes under the scope of *allegedly*, the result is "he shot her and it is said that this happened accidentally"; if the whole predicate of the reported clause comes under its scope, the result is "it is said that he shot her dead and that this happened accidentally".

What *allegedly* and *accidentally* have in common is that they entail an extra perspective layer from which the content they relate to is viewed, thereby modifying the way in which this content should be integrated into the emerging story. In the narrative spaces framework (as in general mental space theory), they would have the role of *space builders*. More precisely, they are linguistic cues that not just prompt any new space, but they prompt a *structured* space or *frame*, a space with a characteristic conceptual structure or topology. In Section

4.3.6 I will discuss this further, focussing on how readers assimilate the topology with local contextual details and integrate the result into the emerging story. Before that, two other issues will be looked at: potential context-specificity of *alleged(ly)* and ways in which compositional and holistic constructions of complexity are used in tandem to represent the entire thoughtscape underlying the Pistorius case in a press release published on the day after the shooting.

4.3.4 Alleged(ly): hedging or coordination of viewpoints?

It should be noted that *alleged(ly)* is commonly added in journalistic discourse related to juridical issues and criminal offence: news media can use it as a way of avoiding responsibility for anything that is still under consideration in court, where they would otherwise be liable to charges of slander or libel. This prompts the question whether *allegedly* is a conventionalised "hedge" associated with criminal and juridical reporting, rather than a more flexibly applicable way of coordinating viewpoints.

It seems possible that a viewpoint package becomes specialised for a particular context, thereby gradually losing the option of being "upacked" and worked out in terms of viewpoint layers, hence of coordinating viewpoints in a more flexible way. Vandelanotte's (2007) corpus research indeed suggests that the adjective *alleged* shows a degree of specialisation, reflected in frequent collocations with words referring to criminal offences and police investigations. In such contexts it normally "transfers" responsibility for a particular (phrasing of a) claim to an authority advancing the charge or dealing with the investigations in a legal case. The example given in Vandelanotte's discussion (2007: 363) illustrates this particular use:

> (4) In New Jersey today, a jury will hear closing arguments in the trial of four young men accused of raping a mentally retarded woman. The **alleged** rape took place in the suburban town of Glen Ridge, New Jersey, four years ago this week.

Assessment of 100 randomly chosen instances from the British National Corpus (BNC) showed that the adverb *allegedly* is also commonly used in this way, but

by no means exclusively.[60] Ample instances (in the current sample at least 20) can be found where the external source of viewpoint referred to by *allegedly* is not some sort of legal authority, but, for example, a particular tradition or history (e.g. "Simon Peter [...] on whom Jesus allegedly founds his church"; BNC, EDY), a source of unconfirmed or even questionable authority (e.g. "one tries to sell [...] smear in a bottle allegedly from the great Madonna herself"; BNC, CBC), or an entity specified in the direct context (e.g. the protesters in "the protesters' vociferously expressed and allegedly 'sincere' ideals"; BNC, HTP); it is unclear from my sample whether the use of quotation marks is indicative of this type of use).

Although an important effect of the insertion of *allegedly* in the press release cited in (3) is the relegation of responsibility for the content under its scope to an external party, I suggest to see this viewpoint package here not primarily as a conventionalised hedge bound to a particular context. Rather, I argue that it is a lexical item cuing the coordination of two viewpoint layers in a way not very different from the working of complementation and inquit-constructions. There are differences in what is in general left implicit or realised "on stage". I expect that this reflects a more general pattern: in the case of complementation or inquit-constructions the party responsible for the external viewpoint appears to be given by default (although impersonal constructions are possible: "It is claimed that..."); in the case of viewpoint packages this party can more easily be left implicit or "off stage" (although it can be elaborated in the context: see the BNC examples cited above). Future (corpus) research will have to shed more light on such differences in emphasis, focussing in particular on how responsibility is distributed over the speaker and the external source of viewpoint, and the degree to which both are on or off stage (see also Langacker, 1990; Wierzbicka, 2006; Vandelanotte, 2009, 2012; Dancygier, 2012a).

[60] Using SketchEngine (http://www.sketchengine.co.uk; see Killgarif et al., 2014) a random sample of 100 instances was drawn from the 1039 instances of *allegedly* in the British National Corpus (BNC; 112,181,015 tokens in total). The BNC is distributed by Oxford University Computing Services on behalf of the BNC Consortium. All rights in the texts cited are reserved. For information, licensing conditions, and use of the text identifiers (the three-letter codes EDY, CBC, and HTP in my citations) see http://www.natcorp.ox.ac.uk.

4.3.5 Representing the thoughtscape

When a larger excerpt of the press release starting with sentence (3) is considered, it becomes clear that multiple means of viewpoint coordination are being combined and mutually embedded to allow readers to form an understanding of the full complexity. Consider the second sentence of the press release:

> (5) He had **mistaken** her for a robber, the *Afrikaans daily* reported on its website.
> (SAPA, 'Oscar Pistorius shoots girlfriend: report', 14 February 2013)

In this sentence, *mistaken* functions as a viewpoint package picking up on the earlier *accidentally*: it entails an extra viewpoint layer by implying that at a certain point the athlete *believed* that something was the case, whereas in actuality it was not. Another inquit-construction can be found: "the *Afrikaans daily* reported [...]" has scope over the clause containing *mistaken*. The past perfect "had mistaken" (as opposed to simple past "mistook") is also involved in viewpoint coordination and has an effect similar to that of *allegedly* in (3): the responsibility for the reported mindstate of the athlete is not (fully) attributed to the *Afrikaans daily* but relegated to a third party. So, example (5) contains a combination of three different means of viewpoint coordination: sentence embedding (here an inquit-construction), usage of the tense system (here past perfect), and usage of a viewpoint package (*mistaken*). However, as announced at the beginning of this chapter, here I focus on a comparison of viewpoint coordination effected by viewpoint packages (holistic) on the one hand and grammatical patterns such as complementation or inquit-constructions

(compositional) on the other, and refrain from detailed analysis of the viewpoint effects of tense, modality, passive voice, or negation.[61]

The next sentence of the press release exhibits not only a combination of viewpoint packaging and embedding, but a form that may be called a hybrid of the two:

> (6) Police spokeswoman Captain Sarah Mcira **confirmed** she [Reeva Steenkamp, MvD] was shot in the arm and head.
> (SAPA, 'Oscar Pistorius shoots girlfriend: report', 14 February 2013)

As a whole, (6) has a form not very different from (3) and (5): it features a reported clause, this time in a complementation construction, attributed to the viewpoint of a specified source. However, a different verb is used: *to confirm* rather than *to report*. Whereas *to report* does not automatically imply information about perspectives other than that of its subject, *to confirm* entails the assumption that the content of the reported clause it introduces was already claimed to be true by someone else; or possibly: that this content is common knowledge. More precisely: for a speaker A to *confirm* proposition p for addressee B, means that A states p and implies that p has been claimed before, and that it is accessible to B that p has been claimed before. In terms of Clark's (1996) view that will be detailed in Chapter 5: p must be part of the *common ground* of A and B. Viewed this way, *to confirm* can be said to be a lexical unit with a topology containing an extra viewpoint layer from which the related content is viewed, and therefore qualifies as a viewpoint package. As a result, the complex of viewpoints involved in (6) can be depicted as follows:

[61] Throughout various traditions of linguistic research, attention has been drawn to the way in which a range of linguistic phenomena are involved in managing viewpoint. For a volume discussing a wide variety of linguistic means for viewpoint coordination, see Dancygier and Sweetser (2012). More examples, among many others, are the description of the disjunctive function of particular adverbs (*allegedly* would come under this category) in Functional Grammar (see Pinkster, 1990: 4, 32ff; Greenbaum, 1969). Wierzbicka (2006, esp. chapter 7 and 8) usefully analyses semantic distinctions between a number of reporting verbs and "epistemic adverbs" (also including *allegedly*), in the context of the supposed English rationalist cultural norm of epistemic caution rather than that of the construction of complex thoughtscapes. For an account of perspectives and the aspectual adverbs *still* and *yet*, see Ter Meulen, 2003.

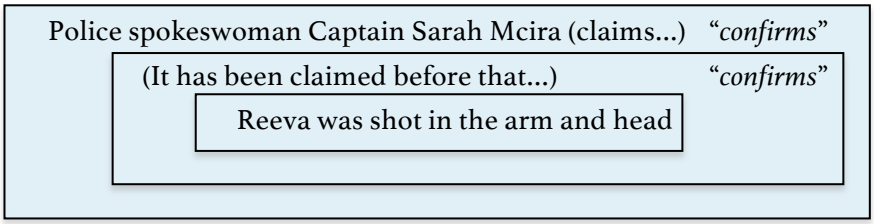

Figure 2

To round off this part, I will take the discussion of how complexity is constructed in the news reports back to Dancygier's narrative spaces framework.

4.3.6. Viewpoint packages, the emerging story, and (de)compression

In the Eggers example cited in (2) there were various linguistic cues signalling to the reader that the empty house, the blood on the walls, the trial, etcetera, were not as such to be integrated into the emerging story. All these crime-related details were embedded in viewpoint layers of Dave's imagination. Instead of concluding that Dave's brother is in severe danger, or perhaps even already dead, readers integrate into the emerging story a more general conclusion: that Dave has hysterical tendencies. This hysteria could be considered a holistic (or *compressed*; see Dancygier, 2012: 100-102) version of the viewpoint layers spelled out in the text, which is not mentioned as such but left implicit for the reader to construe. As we have seen, in some parts of the Pistorius reports we find the precise opposite: here, several of the constituent viewpoint layers are not spelled out in the texts, whereas their holistic counterparts are mentioned: the viewpoint packages *allegedly*, *accidentally*, and *mistaken*. This gives rise to the question whether readers at any point "unpack" (or *decompress*) the packaged structure in their minds; or in other words: whether they in some way construe the separate viewpoint layers contained in the packages.

Part of this question can be answered by looking at the extent to which news items following up on the first headlines build on the packaged viewpoint

layers. Consider the following two excerpts from a news article that was published one week after the shooting:

> (7) Mr. Pistorius told the court [...] he heard a strange noise coming from inside his bathroom, climbed out of bed, grabbed his 9-millimeter pistol, hobbled on his stumps to the door and fired four shots.
> (International Herald Tribune, 'Testosterone reported at home of Pistorius', 21 February 2013)

> (8) Prosecutors said [...] that Mr. Pistorius was calm and had the presence of mind to strap on his prosthetic legs, walk to the bathroom door and open fire as Ms. Steenkamp cowered behind it.
> (idem)

These two possible versions of what could have happened before Pistorius fired his gun present details that are intended to take up either on the scenario of an *accident*, in which the athlete *thinks* he is shooting at a burglar ("strange noise", "hobbled on his stumps"), or on that of *murder*, in which he *knows* he is shooting at his girlfriend ("presence of mind to strap on his prosthetic legs", "walk", Steenkamp "cowered" behind the door). Unless readers have some representation in their minds of these two competing scenarios and the mindstates appropriate to each, it hardly makes sense for newspapers to provide such details without elaborating on their relevance. It may be noted that there is indeed no need for newspapers to do this: details as cited in (8) and (9) can be (and are indeed widely) published without explicit attempts to place them in the story as a whole.[62]

More examples can be found; for instance: by presenting evidence that an argument took place just before the shooting, the prosecutors imply that the

[62] Note that the reports are highly redundant across various newspapers and other media. A week after the shooting, several pieces of information had been added to the "canon" represented in virtually all reports, including witnesses who claimed to have seen light or heard voices in the Pistorius home, small fragments of Pistorius' account of the events to the court, and small fragments of what the police detective leading the investigations has reported (e.g. that he saw a substance which could have been testosterone in the athlete's bedroom). They are often followed by a brief summary of news coverage of the case so far, but no explicit explanations of how they fit together are provided.

athlete *could have known* that Reeva was inside the bathroom when firing the gun:

(9) A witness heard "non-stop shouting" in the home of South African athletics star Oscar Pistorius shortly before his girlfriend was shot dead, the detective leading the murder investigation said on Wednesday.
(Reuters News, 'Witness heard 'non-stop shouting' before Pistorius shooting', 20 February 2013)

(10) While Mr. Pistorius had said the house was dark, the prosecution cited a witness as saying a light had been switched on when the first of four shots was fired. The witness heard a gunshot, then the sound of a woman screaming, then more shots.
(International Herald Tribune, 'Testosterone reported at home of Pistorius', 21 February 2013)

These details are then again challenged by the athlete and his lawyer: the untrustworthiness of these testimonies makes it impossible for the prosecutors *to know* whether or not Pistorius *could have known* that Reeva was in the bathroom:

(11) But the defense disputed that testimony, saying the neighbor who claimed to have overheard an argument in Mr. Pistorius' home in fact lived 600 yards, or about 550 meters, away.
(idem)

In order to be able to integrate such details as presented in (7)-(11) into the emerging story in a meaningful way, readers must already have a particularly structured representation of the thoughtscape in mind. At the very minimum, this representation must comprise the distinction between the two competing scenarios (accident or murder) and several of the viewpoint layers mediating

between the reader and the "actual past event" of the shooting.[63] A schematic depiction can be drawn as follows:

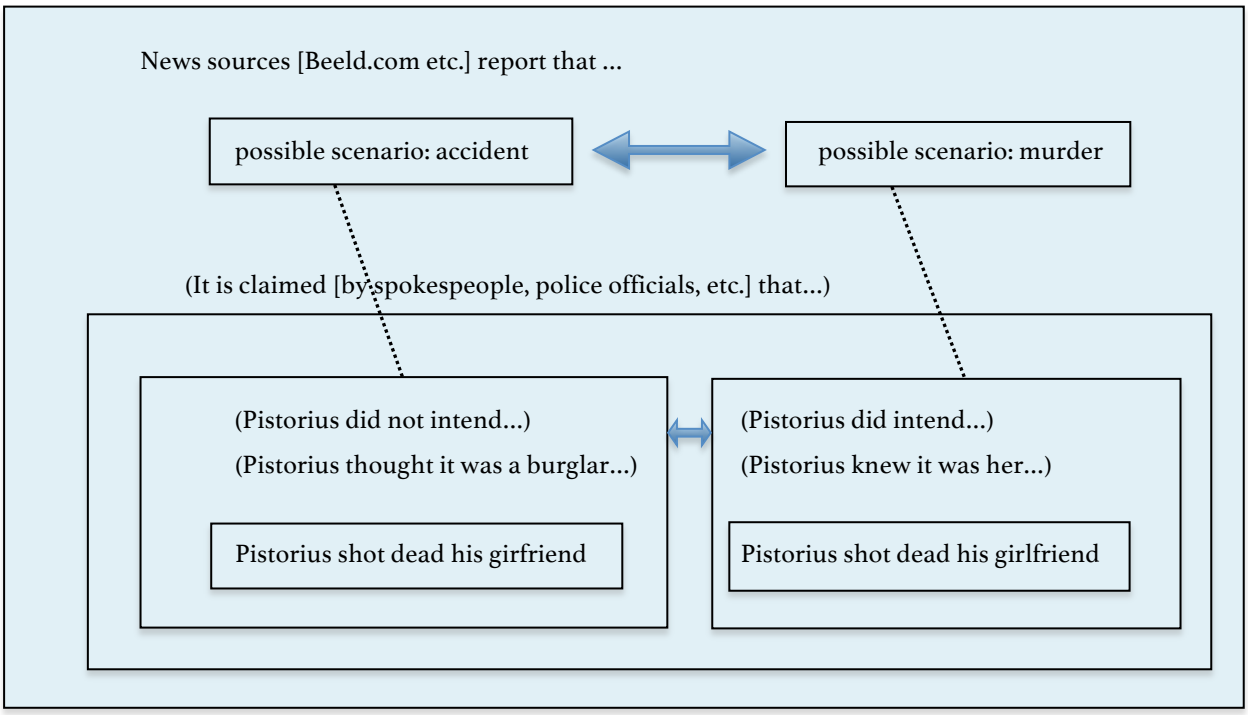

Figure 3

Figure 3 illustrates once more that the claim that Pistorius killed his girlfriend is embedded in an array of viewpoint layers. As stated above, these layers can partly be understood as the result of newspapers relegating responsibility to external sources, such as authorities and other media (see Section 4.3.3 and 4.3.4). However, another reason for their presence in this particular case is that the *actual past event* cannot be retrieved: no one, except Pistorius, knows what "really" happened. As a consequence, reports appearing on the first day after the incident are confined to presenting *possible views* on what has led to the known outcome: that Reeva Steenkamp was shot dead. This is precisely what *accidentally* and *mistaken* do in (3) and (5): they prompt the viewpoint layers of

[63] Cf. Dancygier's (2012) analysis of the opening passage of Margaret Atwood's *The Blind Assassin* and her Figure 4.1 in particular (89-91). There is an important parallel (besides several differences) between the Pistorius case and Dancygier's example: both start out with two alternative scenarios of "what happened", expressed in various viewpoint layers which provide the basis for the elaborations that follow.

Pistorius *thinking* it was a burglar and hence *not intending* to shoot Steenkamp, and arguably at the same time their alternatives of him *knowing* she was in the bathroom and hence *intending* to shoot her—after all, it is widely accepted that a negated proposition also entails its positive counterpart (see Fauconnier, 1997; Sweetser, 2006; Dancygier, 2012a).[64] In this way, the reports and press releases appearing on the first day after the shooting provide what Dancygier calls a "proleptic 'summary'" (2012: 90): they prompt a particular configuration of narrative spaces that persists as a fundament as further elaborations follow.

Regarding the linguistic cues involved in building up this configuration it can be observed in (3), (5), and (6) that sentence embedding is used to coordinate the viewpoints of various sources reporting on the incident. These are the "outer" layers of the schema in Figure 3: the viewpoints of news media and the authorities and their spokespeople. As a result of this choice of linguistic form of viewpoint coordination, the constituting viewpoint layers are presented explicitly in the text (complexity is constructed *compositionally*). By contrast, "inner" layers of the schema, i.e. the representation of the athlete's perspective at the moment of pulling the trigger, completely relies on viewpoint packages until more direct information becomes available several days after the shooting. This pattern is consistent throughout the twenty articles and press releases of the first day after the shooting I have looked at.[65] In general it is thus only on the basis of words such as *accidentally* and *mistaken* that readers can build an understanding of the athlete's possible mindstates when he fired his gun. Given that newspapers publish details following up on precisely these mindstates without further introduction, we may conclude that readers must have some form of access to the packaged (or compressed) viewpoint layers when the context so requires.

[64] This picture is consistent throughout twenty articles and press releases of the first day after the shooting I have looked at: fifteen refer to the shooting with *mistaken*, *mistook*, *mistaking*, *a mistake*, *accidental*, or *accidentally*. Three (of those fifteen) also speculate on the possibility of murder explicitly; the remaining five mention neither of the two alternative scenarios in explicit terms, but arguably imply both. The reports published on 14 February 2013 were retrieved on 21 February 2013 using Lexis Nexis (http://academic.lexisnexis.nl).

[65] The only exception being *Asian News International*: "Pistorius [...] has allegedly shot his girlfriend to death in the early hours of Thursday morning after thinking she was an intruder, reports claim".

Put differently: it seems that readers do not unpack (or decompress) packages by default—a package may usually be closed and taken on board holistically—but when necessary, the implied viewpoint layers can be accessed. This leads to a final core aspect of viewpoint packages. Consider the word *mistaken*: thanks to its holistic nature it does not automatically necessitate reflecting on a whole series of assumptions about "who knew what at which moment in time". Still, it is easy to prompt assumptions about the distribution of such information over minds by providing specific contextual details along with it. This has clear communicative, and conceivably also cognitive, advantages: constructing the viewpoint complexity underlying a situation compositionally every time it needs referencing (e.g. by using sentence embedding) is possible, but often unnecessarily ponderous. Therefore, in contexts where space and time are limited (as is the case in news reports and headlines), holistic introduction of complexity using viewpoint packages appears to be a preferable option. The difference between the actual newspaper quotes in (3), (5), and (6) and the unpacked thoughtscapes as explicated in the introduction of this chapter ("news media *suggest* that various sources *report* that the athlete *claimed*…etc.") and depicted in Figure 1-3 testifies to that: whereas the first will be perceived as natural and everyday formulations, the latter present the same information in the form of a layered structure that cannot be absorbed at a glance.

4.4 Conclusion and discussion

4.4.1 Future directions: processing and acquisition

In the past two decades, a fair amount of attention in cognitive and evolutionary psychology has been devoted to the cognitive challenges humans face in their social environments. In order to function well socially it is of great importance to be able to reason about mindstates of others, that is: assess what they think, believe, intend, desire, and so on. As a result, primate social life is

highly demanding of cognitive resources.[66] Dunbar (e.g. 2003) suggests that humans typically form social networks of around 150 individuals. In his view, to maintain such a network, humans must be able to process what has been termed *multiple-order intentionality* up to five or six orders. Examples of such assumed processing tasks involving four orders (12) and five orders (13) are:

(12) Jenny *hoped* the greengrocer *believed* the chemist *had wanted* to give Emma a job. (Stiller & Dunbar, 2007)

(13) (a. Connie *knew* that John *suspected* that Pete *thought* that Sheila *hoped* that John would ask her out. (idem)

Note that these statements show similarities to the phrasing of the thoughtscape underlying the Pistorius case in Section 4.1 of this chapter and to the layered structure depicted in Figure 1-3. Various studies have used such statements in multiple-choice questions to test a participant's abilities to process multiple-order intentionality (see Chapter 6). Where participants generally make few mistakes in questions covering up to fourth-order intentionality, error rates increase quite drastically in questions that involve fifth- or sixth-order. This has led to the suggestion that humans face a natural limit at this point (Kinderman, Dunbar, and Bentall, 1998; Launay et al., 2015).

However, in Chapter 2 I have argued that the way in which situations involving multiple-order intentionality (which I call thoughtscapes) are represented in discourse, greatly influences the actual performance by humans in processing such tasks. Consider Shakespearean drama: by the end of act II, the audience of *Othello* has to *understand* that Iago *intends* that Cassio *believes* that Desdemona *intends* that Othello *believes* that Cassio *did not intend* to disturb the peace. When represented like this, using an embedded sentence structure, this is highly opaque, whereas it is beyond doubt that the play has been

[66] This is true throughout the entire primate world, where social groups are structured around dyads: every individual has a personal relationship to some or all of the other individuals in the group. Having a personal relationship with someone, involves keeping track of a lot of knowledge about this individual, including what this individual knows about other group members and oneself (e.g. David-Barrett & Dunbar, 2013). This idea is central in work on the *social brain hypothesis* (e.g. Dunbar, 2003; Byrne & Whiten, 1988). See Chapter 1, Section 1.3 for more details.

understood and appreciated by innumerable different audiences for several centuries. I have suggested that in *Othello* and comparable stories, certain "expository strategies" characteristic to narrative facilitate the effective representation of complex networks of mindstates, thereby alleviating the reader's or spectator's burden of processing such complexity cognitively. In Chapter 3 I have come to a similar conclusion when analysing excerpts from several novels.

In this chapter I have suggested that viewpoint packages serve to communicate complex thoughtscapes efficiently and naturally. They allow readers to take on board certain parts of a layered structure holistically, while the underlying complexity can be unpacked (or decompressed) if the context so requires, but need not be otherwise. Extending this point to the domain of cognitive processing, I suggest that, in psychological terms, viewpoint packages have *Gestalt*-like properties: they can be used as holistic items, while their constituent components remain accessible (Lakoff, 1977; see also Gigerenzer, Hertwig, & Pachur, 2011; Humphrey, 1924). As such, viewpoint packages could provide crucial scaffolding for the cognitive handling of complex thoughtscapes: given constraints of working-memory and processing, working with packages allows humans to engage in more complex patterns of reasoning than working with substantively identical, but non-packaged primitives (Beekhuizen and Van Duijn, 2013). This is an alley, I propose, that should be further explored in the future. It would fit with the broader idea that human cognition relies crucially on "tools" obtained through social learning and potentially accumulated through cultural evolution (see Tomasello, 1999; Dennett, 2000). Viewpoint packages can be considered such tools: when we learn to use them in language, their representational power opens new worlds of possibilities for our mental and communicative activities, not replacing but complementing existing ones.

Finally, it is worth adding that there is another interesting connection between viewpoint packages and complementation constructions within the realm of language acquisition. Diessel & Tomasello (2001) have demonstrated that the earliest items that *look like* instances of complement taking predicates (thus of complex syntax) in children's utterances actually have the status of formulaic items that mark subjectivity or illocutionary force (*I think …*, *I wish …*,

Look ...). Thus, they may at that stage be regarded as a kind of viewpoint packages in the sense of this paper. In a developmental perspective, they differ from the other viewpoint packages discussed here in that they are relatively simple (but this is a matter of degree), and that they form the starting point of the growth, in due course of development, of a network of complementation constructions that also includes productive abstract patterns (cf. Verhagen, 2005: 110). The latter is not the case for the English adverb *allegedly*, but the process of the acquisition of such items and the social cognitive capacities underlying it (both in initial and in later acquisition) may be assumed to be similar, if not the same. I leave it to future research to elaborate and test this intriguing possibility.

4.4.2 The refined definition of viewpoint packages

Throughout this chapter I have looked at newspapers reporting on the Pistorius case, finding that already on the first day after the shooting a complex thoughtscape was covered by headlines and news articles. This complexity (depicted schematically in Figure 3) was partly cued compositionally, with the constituting layers in the text, and partly holistically, using words such as *alleged(ly)*, *accident(ally)*, and *mistake(n)*. I have proposed to label such words "viewpoint packages" and provided a preliminary definition, which can now be refined: a viewpoint package is a single lexical item that entails at least one implicit viewpoint layer. This (or these) layer(s) allow for the attribution of (parts of) the content to the viewpoint of particular discourse participants that can remain unspecified (as was the case with *allegedly*) or identified in the context (for example "the athlete" with *accidentally*). In this way, a viewpoint package can be used to regulate the (perception of the) amount of responsibility that is taken by the speaker and other involved parties exhibiting intentionality.

In addition, a viewpoint package lets language users take on board complexity holistically, while underlying layers remain accessible: it can be decompressed if the context so requires. This makes it an efficient tool for communicating complex thoughtscapes, and possibly also for processing them cognitively.