Linguistics and Philosophy (2023) 46:1441–1465
https://doi.org/10.1007/s10988-023-09392-3

ORIGINAL RESEARCH

Truth and directness in pictorial assertion

Lukas Lewerentz[1] · **Emanuel Viebahn**[2]

Accepted: 15 May 2023 / Published online: 14 July 2023

Abstract
This paper develops an account of accuracy and truth in pictorial assertion. It argues that there are two ways in which pictorial assertions can be indirect: with respect to their content and with respect to their target. This twofold indirectness explains how accurate, unedited pictures can be used to make false pictorial assertions. It captures the fishiness of true pictorial assertions involving target-indirectness, such as true pictorial assertions involving outdated pictures. And it raises the question whether target-indirectness may also arise in linguistic assertion.

Keywords Assertion · Pictorial assertion · Pictorial accuracy · Pictorial content · Literalness · Super linguistics

1 Introduction

Since the rise of mobile messaging and the smartphone, a substantial and growing portion of everyday communication is based on pictures. Indeed, in many situations pictorial communication is nowadays as natural as linguistic communication. Consider the following example:

> *Eiffel Tower 1*. Sybil has just started a new job and a friend messages her to ask whether the Eiffel Tower is visible from her desk. The Eiffel Tower is indeed

✉ Emanuel Viebahn
emanuel.viebahn@fu-berlin.de

Lukas Lewerentz
lukas.lewerentz@phil.uni-giessen.de

[1] Institut für Philosophie, Justus-Liebig-Universität Gießen, Rathenaustraße 8, D-35394 Gießen, Germany

[2] Institut für Philosophie, Freie Universität Berlin, Habelschwerdter Allee 30, D-14195 Berlin, Germany

> visible from Sybil's desk. Sybil uses her smartphone to take a picture that shows her desk in the foreground and the Eiffel Tower outside the window and sends it to her friend (with no further comment).

In this situation, it is just as natural to send the picture as it is to respond by writing "Yes, it is visible from my desk." Furthermore, these two communicative acts appear to be very similar. Through both acts, Sybil puts forward that the Eiffel Tower is visible from her desk. And the two acts seem to have the same illocutionary force: in both cases, it seems, Sybil *asserts* the content in question. This pre-theoretical judgement fits with philosophical views that allow for pictorial assertion, which have been advanced e.g. by Kjørup (1974), Eaton (1980), Wolterstorff (1980, 201), Walton (1983, 79), Korsmeyer (1985), Viebahn (2019), Dixon (2020), Kulvicki (2020, 1) and García-Carpintero (2023).

But despite the ubiquity of pictorial communication and the fact that many philosophers accept that pictures can be used to make assertions, the *mechanisms* of pictorial assertion are underexplored. This paper is an attempt to start to fill this gap by elucidating the relation between pictorial accuracy and the truth of pictorial assertion. The conditions under which *pictures* are accurate have received quite some attention in the literature. But pictorial accuracy and the truth of pictorial assertions are separate matters, as the following variations on the first example illustrate:

> *Eiffel Tower 2*. The Eiffel Tower is in fact *not* visible from Sybil's desk. To respond to her friend's question, Sybil takes a photo of the view from her desk, edits in a photo of the Eiffel Tower, and sends the edited photo to her friend (with no further comment).

> *Eiffel Tower 3*. The Eiffel Tower is in fact *not* visible from Sybil's desk. To respond to her friend's question, Sybil takes a photo of the view from her colleague's desk, from where the Eiffel Tower is visible, and sends it to her friend (with no further comment).

In both variations, Sybil seems to be making a false assertion. But while the photo in the second case is edited and thus inaccurate, there is an intuitive sense in which the photo in the third case is perfectly accurate. In the latter case, it seems, Sybil uses an accurate picture to make a false assertion. Situations of this kind raise puzzling questions: How can accurate pictures be used to make false pictorial assertions? How can the difference between the two communicative acts be adequately captured? And, more generally, how are the accuracy of pictures and the truth of pictorial assertions related?

Despite the growing interest in the semantics and pragmatics of non-linguistic communication (see Schlenker, 2018; Abusch, 2020; Schlöder & Altshuler, 2023 and other papers in the novel field of *super linguistics*), questions of this kind have rarely been discussed.[1] Our aim in this paper is to shed light on them by developing

[1] To our knowledge, Novitz (1975) and Korsmeyer (1985) are the only papers offering some brief discussion of the matter. Greenberg (2018) discusses "assertoric uses of pictures" (p. 871), but he is concerned

an account of accuracy and truth in pictorial assertion. A key role in this account is played by the tenet that there are two ways in which pictorial assertions can be indirect: with respect to their content and with respect to their target. We will try to show that the difference between *Eiffel Tower 2* and *Eiffel Tower 3* should be captured in terms of target-indirectness, and we will defend our approach against an alternative approach that denies the possibility of target-indirectness. Furthermore, we will argue that directness matters: it matters for pictorial communication, and in particular for assertions made with photos, as the photos presented are commonly used to provide evidence for the content asserted. And it raises the question whether target-indirectness may also arise in linguistic assertion.

Our focus in what follows will be on photographic assertions, as these are the clearest cases and most common among pictorial communicative acts. But what holds for photographic assertions should hold for at least some pictorial communicative acts that do not involve photos, e.g. for communicative acts made with drawings, outline tracings and carbon copies.[2] Moreover, once pictorial communicative acts are in view, one might wonder whether yet other acts should be counted as assertions, e.g. acts of showing (directly, without a representational device).[3] We are open to that possibility, but will here restrict our attention to communicative acts involving representational devices, such as words or pictures, and (as mentioned) will focus mainly on photographic communicative acts.

2 Pictorial assertion

To begin with, we would like to provide some support for the possibility of pictorial assertion. While it is exceedingly natural to treat Sybil's pictorial communicative acts as acts of assertion, some theorists will be bound to disagree on this point. In particular, views that entail that asserted content must be the semantic content of linguistic expressions uttered rule out the possibility of pictorial assertion (see e.g. Carson, 2010, 18 and Stokke, 2018, 31). But there are good reasons to hold that pictorial assertion is possible, as becomes evident if we compare Sybil's act of sending the photo with her written response "Yes, it is visible from my desk," which clearly is an assertion.

Firstly, both communicative acts fulfil central criteria of the main theories of assertion: whether Sybil sends words or a picture, she (i) presents herself as believing that the Eiffel Tower is visible from her desk (Bach & Harnish, 1979), (ii) proposes to add the aforementioned content to the common ground (Stalnaker, 1999), (iii) takes on a commitment to that content (Peirce, 1934; Brandom, 1983; MacFarlane, 2011)

with *pictures* and their accuracy ("my agenda here is to ask […] what makes a picture *accurate* or *inaccurate*" (2018, 869)), and not with how pictures are *used* to make assertions. For a recent illuminating study on pictorial pragmatics (though not with a focus on the accuracy of pictorial assertion), see Esipova (2021).

[2] We discuss outline tracings and carbon copies in footnote 29.

[3] Relatedly, see De Leon (2022) for arguments that pointing gestures can be used to "*say something* in a conversation" (p. 2). De Leon, however, restricts the notion of assertion to linguistic assertion; see pp. 13–14.

and (iv) performs an action that she shouldn't perform if she believes the content put forward to be false (as is entailed by Williamson, 1996). As mentioned, some theories in these broad groups also require assertions to be linguistic: Stokke (2018, 31), whose view of assertion falls into group (ii), holds that assertions are proposals to update the common ground *made by uttering a sentence*. Similarly, Carson (2010, 18) defends a view of assertion that falls into group (iii), but also holds that asserting requires *producing a linguistic token*. However, these requirements for assertion to be linguistic are separate from the central criteria in all four groups; they are add-ons. The criteria themselves are neutral with respect to the communicative medium, and so pictorial assertions are compatible and possibly even predicted by the main theories of assertion. For example, if it is possible to use pictures to propose an update to the common ground, which reasons do theorists in group (ii) have to rule out pictorial assertion? Or, if it is possible to take on assertion-like commitment to a content through presenting a picture, why should this communicative action not count as an assertion for theorists in group (iii)? In the absence of reasons not to allow pictorial assertion, the similarity between Sybil's pictorial act and her linguistic act suggests that pictures can be used to assert.

Secondly, there is (indirect) empirical evidence that ordinary speakers consider communicative acts that are like Sybil's act of sending the picture to be assertions. In a recent study, Viebahn and Wiegmann (2023) presented subjects with vignettes in which agents use either a written message or a picture (without any accompanying text) to put forward a proposition they believe to be false. Subjects were then asked to judge whether the communicative act counts as a lie. The results show that subjects were just as likely to judge the pictorial acts as lies as they were to judge the corresponding linguistic acts as lies, even if they were given the opportunity to classify these acts as misleading (but not lies). These results support a folk view that permits pictorial lies. Lies are standardly considered to form a subset of assertions: it is widely held that lying requires asserting something one believes to be false (see Stokke, 2018 and Viebahn, 2021). So if theories of assertion are meant to be compatible both with the folk view of lying and with the standard assumption of lies as a subset of assertions, this gives us a second reason to accept the possibility of pictorial assertion.

Much more would have to be said to provide a comprehensive case in favour of pictorial assertion. But we do think that these considerations motivate a unified treatment of what we would call linguistic assertions and pictorial assertions. And even if the two phenomena are not grouped under the same label of *assertion*, insights can be gained from treating them together, as we will do in what follows.

3 Pictorial accuracy

What is it for a picture to be accurate or inaccurate? A detailed, recent discussion of this question is due to Greenberg (2018), who argues for a three-part model of pictorial accuracy. The three parts of this model are (i) the singular content of the picture, (ii) the attributive content of the picture, and (iii) the target scene the picture aims at.

The distinction between singular content and attributive content goes back to Goodman (1968). The *singular content* of the picture is best thought of as the subjects that are represented in the picture; it is what a picture is *of*. The *attributive content* includes the properties the picture ascribes to its subjects; it represents the subjects *as being a certain way*. For example, if Sybil draws a sketch of the banana plant in her office, then the singular content of the sketch is the banana plant (plus possibly some of its surrounding items, such as its pot), while the attributive content represents the banana plant as being a certain way, such as having a certain number of leaves of a certain shape and colour. The traditional construal of pictorial accuracy appeals only to a picture's singular and attributive content: a picture is accurate just in case the attributive content holds of the singular content. Goodman (1968, 38) endorsed a view of this kind, stating that for a picture to be accurate is "for the object represented to have the properties that the picture in effect ascribes to it".[4]

The main aim of Greenberg (2018) is to show that such a construal of pictorial accuracy is impoverished. According to Greenberg, pictorial accuracy depends not only on singular content and attributive content, but also on the *target scene* a picture aims at. A picture's target scene is not part of the content of the picture; rather, it is the scene that the picture, with its content, is meant to capture. We can ask of a picture whether its content accurately captures this or that scene, and our answer may not always be the same. A picture depicting the Eiffel Tower as illuminated in red may accurately represent a scene in Paris on some specific night while not accurately representing a scene on a different night. In that case, the picture's content holds of the first target scene (since the Eiffel Tower is in fact illuminated in red on that night), and so the picture is accurate with respect to that target, but the picture is inaccurate with respect to the second target scene, because its content doesn't hold of that scene (since the Eiffel Tower is in fact not illuminated in red on that night).

Importantly, pictures show the situations they aim to depict from a certain perspective, and so whether a picture is accurate or not also depends on what that perspective is. For example, a drawing may aim to depict the scene on the table in front of us from the point of view of the door. That drawing would misrepresent the scene if it depicted the red mug on the table as being to the left of the book that's also on the table – even if there is another scene that the same drawing does capture accurately, for example, the table as seen from the window that is opposite the door. To capture this perspective-dependence of pictorial accuracy, Greenberg suggests that the target scenes that pictures aim to depict incorporate an orientation, a point of view from which the depicted situation is seen. Accordingly, Greenberg construes the target scene of a picture as "a possible spatio-temporal situation anchored at a particular viewpoint" (2018, 870), which he in turn models as a *viewpoint-centred world*. He

[4] Similar accounts of pictorial accuracy are offered by Casati and Varzi (1999, 194–195), Rescorla (2009, 180) and Greenberg (2013, 252). As Greenberg (2018, 867), we will assume that the attributive content of pictures includes high-level properties, such as *having a certain number of leaves* or *being hydrated*. Some theorists have identified levels of pictorial content that are not suitable for such high-level properties, but there is widespread agreement that high-level properties can feature at *some* level of pictorial content. For example, we shouldn't expect to find high-level properties in the *basic representations* of Hyman (2006, 63) or in the *bare bones content* of Kulvicki (2020, 25–27), although Hyman and Kulvicki do allow such content at less basic levels of pictorial representation.

takes a viewpoint to be "an oriented location" (2018, 870), which we might think of as a location with an (invisible) arrow pointing in some direction.

The idea, then, is that a picture is accurate just in case its content (which includes both the picture's singular and its attributive content) holds of its target scene. Here is Greenberg's (2018, 876) model of pictorial accuracy in full:

> *Greenberg's three-part model of pictorial accuracy*
> A picture P is accurate in a context c if and only if the attributive content expressed by P in c is instantiated by the singular content expressed by P in c in the target scene selected by c.

We will work with this model of pictorial accuracy in what follows.[5] As the distinction between attributive and singular content is not of importance for this paper, we will simplify matters by concentrating on the following version of Greenberg's model, which only mentions the content and target of a picture:

> *Greenberg's three-part model of pictorial accuracy simplified*
> A picture P is accurate in a context c if and only if the content expressed by P in c holds in the target scene selected by c.

This should be in Greenberg's spirit, as he notes that "singular and attributive content make up an integrated whole" (2018, 867). Although the simplified model only mentions two aspects of a picture that are relevant for its accuracy, it is still a three-part model, insofar as attributive and singular content together determine the content expressed.[6]

As Greenberg (2018, 874) points out, the distinction between the content of a picture and its target parallels Kaplan's (1989, 522) well-known distinction between the content of an (occurrence of a) sentence and the circumstance of evaluation at which that content is then evaluated for truth. Just as, on Kaplan's view, the content of a sentence may be true at one circumstance of evaluation while being false at another, the content of a painting may be true at one target scene while being false at another. That it makes sense to ask whether the content expressed by some sentence or picture would be true in this or that circumstance (or this or that target scene) leaves open what the *relevant* circumstance (or target) is for evaluating whether a given utterance of that sentence (or use of that picture) is true.

Two further aspects of Greenberg's account will be important and are thus worth noting here. The first concerns *the role of context* for pictorial accuracy. Greenberg holds that the context of creation plays an important role in fixing both the content and the target of a picture. The *content* of a picture depends on the context of creation in a

[5] Blumson (2014, 154) accepts a similar view of pictorial accuracy. He holds that "depictions are accurate when the states of affairs they are of obtain, and inaccurate when the states of affairs they are of fail to obtain." Ross (1997, Chap. 5) also defends the view that the contents of pictures are viewpoint-relative and therefore more fine-grained than sets of possible worlds.

[6] As far as we can see, most of the considerations to follow could also be made with the traditional two-part model of pictorial accuracy in the background. But, as will become apparent, Greenberg's model is particularly helpful in analysing the cases of this paper.

twofold way (2018, 868–869): on the one hand, the context of creation fixes an operative system of depiction, on the other, the creator's intentions (as one aspect of the context of creation) fix the singular content of a picture.[7] The *target* of a picture is also context-dependent, though different aspects of the context determine the target for different kinds of pictures (2018, 872). For non-mechanically produced pictures, such as drawings, the creator's intentions are decisive. For photos and other mechanically produced pictures, it is not intentions, but the function and settings of the picture-taking device that matter. A photo may thus aim at a certain target scene (usually the scene before its lens) even if there are no relevant target-intentions in the context, e.g. if the camera is dropped and the shutter is activated by accident. Though context plays different roles for different aspects of a picture, and though there is even variation with respect to which aspect of context matters for one of these aspects (namely the selection of a target scene), it is always the context *of creation* that matters, according to Greenberg, not the context in which the picture is used or in which it is interpreted.[8]

A second noteworthy point is that Greenberg's model is an account of *perfect* pictorial accuracy (2018, 875). Pictorial accuracy can arguably come in degrees: some pictures are more accurate than others. Greenberg is interested in the uppermost boundary of the scale of accuracy. Only in the absence of *any* misrepresentation is a picture perfectly accurate. We will follow Greenberg in using *accuracy* to pick out the notion of perfect accuracy.

In addition, we think it is useful to introduce the notion of *partial pictorial accuracy*. Pictures are usually informationally rich. As a result, a picture may represent some aspects of its target accurately while misrepresenting others. For example, a picture of a green plant in front of blue curtains may accurately represent the colour of the plant's leaves as being green while misrepresenting the colour of the curtains as being red. In that case, the picture would be accurate with respect to the leaves being green while being inaccurate with respect to the curtains being red. One way to think about this is that pictures express a multitude of propositions, and that each of the propositions expressed may hold or fail to hold of the picture's target scene.[9] A picture is accurate with respect to a proposition p that is expressed by the picture just in case p holds of the picture's target scene:

> *Partial pictorial accuracy*
> A picture P is accurate in a context c with respect to a proposition p if and only if the proposition p is expressed by P in c, and p holds of the target scene selected by c.

[7] A system of depiction is "the pictorial analogue of a language" (Greenberg, 2018, 868). For example, colour photos and black and white photos differ with respect to their system of depiction: in a colour photo, a grey patch represents something as being grey, which is not the case in a black and white photo.

[8] The view that it is the context of creation that matters for determining a picture's content is widely shared, see e.g. Lopes (1996), Hopkins (1998) and Abell (2009). Greenberg (2018, 873, n. 13) briefly discusses the possibility of *repurposing* a picture to aim at a new target. We consider this possibility below.

[9] This is not to deny that the contents expressed by a picture also form an integrated whole. One could provide an account of complete pictorial accuracy based on partial accuracy as follows: a picture P is completely accurate in c if and only if P is accurate in c with respect to every proposition p expressed by P in c.

In the example, the photo expresses the proposition that the leaves are green but also the proposition that the curtains are red. It is thus accurate with respect to the former proposition but inaccurate with respect to the latter. When we evaluate a picture for accuracy, we may sometimes focus on some of the propositions expressed while ignoring others, a point to which we'll return below.[10]

4 Truth and accuracy in pictorial assertion

In the previous section, we introduced an account of the conditions under which pictures are accurate. Against this background, we now want to move on to discuss how pictures can be used to make assertions and under which conditions such pictorial assertions are true. Our aim is to bring out two interestingly different ways in which pictorial assertions can deviate from the pictures used to make them: by having either a different content or a different target. That such deviation (or indirectness, as we'll call it) is possible is one important reason why one should distinguish between the accuracy of pictures and the truth of pictorial assertions.

4.1 Assertoric truth

Assertions are often assessed for truth. Some assertions are true, others are false. If someone asserts that London is in England, the assertion is true because the asserted content is true. If someone asserts that London is in Scotland, the assertion is false because the asserted content is false.

Just like pictures, assertions also have targets: roughly, the target of an assertion is where the asserted content has to be true in order for the assertion to be true. That assertions have targets seems to us to be a widely held view, even though it is controversial what exactly these targets are.[11] On one view, the target of an assertion is a possible world, and the assertion of a proposition p is true just in case p is true at that possible world. In this case, the target of an assertion would plausibly be the world at which the assertion takes place: an assertion that London is in England is true if made in our world but would be false if made in a possible world where London is in Scotland. The view that the targets of assertions are possible worlds goes well with the view that the content of an assertion is a proposition and that propositions are true or false at possible worlds. But that the target of an assertion is a possible world is just one view; other views may take targets to be pairs of a world and a time, or situations,

[10] For related comments, see Korsmeyer (1985, 260–261), Viebahn (2019, 248), and Kulvicki (2020, 34–35, 51–52). See also Kulvicki (2020, Chap. 3) for discussion of why a picture can express multiple propositions rather than always just one. Very roughly, the idea is that parts of a picture represent things just like the whole picture does, and that different parts (or more precisely: what they represent) may play different roles when the picture is used in communication – sometimes, some parts may just play no role at all. See also footnote 13 below.

[11] For instance, any view that holds (1) that an assertion is true (false) just in case it has a true (false) proposition as content, and (2) that propositions are true (or false) at circumstances of evaluation (whatever these are; for the notion of a circumstance of evaluation, see Kaplan, 1989), seems committed to the view that assertions have such circumstances as targets. See e.g. MacFarlane (2014, 126–127), for a brief discussion of what he calls "accuracy" of speech acts. Also see Gariazzo (2019) for relevant discussion.

or viewpoint-centred worlds, or something else; and nothing said so far excludes the possibility that different kinds of assertions have different types of targets.

Leaving the nature of targets open, and leaving open how the target of an assertion is determined, we can nonetheless state the following principle:

Truth of Assertions
An assertion of a content p, aimed at a target a, is true if and only if p is true at a.

If all assertions can be assessed for truth, this also holds for assertions made with pictures: assertions we make by using pictures are true or false even if we hesitate to count the pictures themselves as true or false:

Truth of Pictorial Assertions
A pictorial assertion of a content p, aimed at a target a, is true if and only if p is true at a.

On Greenberg's view, the content of a picture holds or fails to hold at target scenes, which Greenberg takes to be viewpoint-centred worlds. It seems plausible that the content of a pictorial assertion has the same kind of content as the picture used to make the assertion. If that's correct, the content of a pictorial assertion also holds or fails to hold at target scenes: the targets of pictorial assertions are also target scenes.

4.2 Directness

On the view just sketched, assertions have contents and targets, and so do pictures. One might therefore wonder how the content and target of a picture are related to the content and target of a pictorial assertion that is made with the picture. Often, they will be closely related. Imagine that a friend texts Sybil, who is still in her office in Paris, and asks whether there are clouds above the Eiffel Tower. In response, Sybil takes a picture which shows the Eiffel Tower and the cloudy sky above it, and she sends it to her friend. Plausibly, Sybil has thereby asserted that there are clouds above the Eiffel Tower. The picture she used is accurate with respect to the proposition that there are clouds above the Eiffel Tower. The proposition is part of the content of the picture, and it holds of the picture's target scene. Moreover, the picture's target scene is plausibly the same as the assertion's target: it is the world as seen from Sybil's desk through the office window, which happens to face the Eiffel Tower. If the picture's target is the same as the assertion's target, the assertion is guaranteed to be true if the picture is accurate with respect to the asserted proposition. Therefore, not just the contents and targets of the assertion and the picture are closely related, their truth and accuracy are so, too.

Let's say that a pictorial assertion is *direct* just in case the asserted content is expressed by the picture used and the target of the assertion is the same as the picture's target scene. As we have just seen, Sybil's assertion that there are clouds above the Eiffel Tower is direct in this sense. The notion of directness is closely related to the notion of being literal. Roughly, a speech act counts as literal when the sentence used to make the speech act has the same content as the speech act, and it counts

as nonliteral when its content differs from the content of the sentence used.[12] If we wanted to expand this notion to the pictorial case, we could say that a pictorial assertion that is direct is literal.

Pictorial assertions could fail to be direct in two different ways. Firstly, the asserted content could fail to be expressed by the picture. Secondly, the assertion's target could fail to be the same as the picture's. It is therefore useful to distinguish two corresponding types of directness:

Content-directness (c-directness)
A pictorial assertion of a proposition p, aimed at target a, made with picture P, is direct with respect to its content (*c-direct*) if and only if p is expressed by P.

Target-directness (t-directness)
A pictorial assertion of a proposition p, aimed at target a, made with picture P, is direct with respect to its target (*t-direct*) if and only if a is the target of P.

A pictorial assertion is thus direct just in case it is both c-direct and t-direct. Moreover, the two types of directness are in principle independent; a pictorial assertion could be c-direct without being t-direct, and vice versa. In what follows, we want to explore what role these different ways of being direct or indirect play for pictorial assertion.

It may be worth highlighting that content-directness, as we have defined it, does not require that the assertion's content encompasses the whole content of the picture used. All that matters for a pictorial assertion to be c-direct is that the asserted content is part of the content of the picture. The content of Sybil's assertion that there are clouds above the Eiffel Tower does not exhaust the content of the photo she uses; for example, it is also part of the content of the photo that the clouds are grey, have a certain shape, and so on. Nevertheless, Sybil's pictorial assertion doesn't strike us as indirect in any substantial sense. Content-directness captures this impression.

That the asserted content doesn't always exhaust the picture's content means that a (direct) pictorial assertion can be true even if the picture used is only partially accurate. For example, Sybil's pictorial assertion that there are clouds above the Eiffel Tower would be true even if the clouds above the Eiffel Tower didn't have the exact shape depicted in the picture she sends to her friend. That the asserted content can be partial also raises the question of what determines which part of a picture's content is asserted and which part isn't. One possible view is that the selection is a matter of the asserter's communicative intentions; on this view, that Sybil (merely) asserts that there are clouds above the Eiffel Tower is explained by the fact that this is what she intends to assert. Another view is that the selection is determined by the question the assertion addresses; on this view, that Sybil (merely) asserts that there are clouds above the Eiffel Tower is explained by the fact that Sybil's assertion is an answer to the question whether there are clouds above the Eiffel Tower, to which other parts of

[12] This notion of literalness does not take into account sameness or difference of target. We'll return to this issue below.

the picture's content are not relevant.[13] It would be interesting to spell out and compare these (and perhaps other) views, but for our purposes, it's not crucial which of them is correct, and so we won't argue for a particular view here. Whatever the exact mechanism is, it seems plausible that a (direct) pictorial assertion's content need not be the picture's full content.

4.3 Indirect pictorial assertions

For all we have said so far, all pictorial assertions are direct. In this section, we want to motivate the idea that not all pictorial assertions are direct. To that end, we will sketch examples that we take to be good candidates for being indirect pictorial assertions.

Let's first consider whether there are pictorial assertions that are c-indirect, that is, assertions whose content is not part of the content of the picture used to make the assertion. To introduce the phenomenon, let us briefly turn to linguistic assertion. It is sometimes said that linguistic assertion has to be "open, explicit and direct" (Pagin & Marsili, 2021). The most straightforward way to spell out this view is as claim that all linguistic assertions are c-direct: in order to (linguistically) assert p, a speaker has to utter a sentence that has p as its semantic content.[14] Such a view rules out the possibility of linguistic assertions where the asserted content and the content of the sentence aren't identical: assertions with non-literal utterances, i.e. utterances featuring metaphor, hyperbole or irony. But it is not clear that metaphorical utterances should be ruled out as assertions. Arguing in support of metaphorical assertion, Bergmann (1982, 231) puts forward the following example:

(1) The nuclear reactor is a time bomb.[15]

Bergmann argues that a speaker can use (1) to assert that the nuclear reactor is likely to fail. We are inclined to agree with Bergmann. If Bergmann is right, linguistic assertion need not be explicit and direct, and non-literal linguistic assertion is possible: a speaker can assert p without uttering a sentence that has p as its semantic content.[16]

[13] Thank you to an anonymous referee for suggesting the question-based view. Relatedly, Esipova (2021) argues that the content of pictures is divided into at-issue and not-at-issue content, just like linguistic content. Esipova (2021, 11–12) also argues that the question under discussion is among the factors that determine (or at least help the interpreter figure out) which parts of a picture's content are at-issue. The connection to Esipova's distinction deserves further investigation. One hypothesis could be that the content that is asserted just is the at-issue content. But note that at least in the linguistic case, speakers are usually committed to not-at-issue content they express. For example, an utterance of "Kim can't apply for this grant again." (Esipova, 2021, 3) usually commits the speaker to the (not-at-issue) proposition that Kim has applied in the past, and the speaker would arguably be lying if they knew that Kim has never applied (cf. Viebahn, 2020). It is not obvious that the same holds for content that is part of a picture used to make a pictorial assertion but isn't asserted. For example, if Sybil uses a photo to assert that there are clouds above the Eiffel Tower, does she commit herself to the proposition that the clouds above the Eiffel Tower are grey and have a certain shape? We are not certain that she does.

[14] A view of this kind is defended by Alston (2000, 120).

[15] We have modified the example slightly to ease discussion.

[16] Among the theorists who have argued for the possibility of indirect assertion are Bach and Harnish (1979, 70 ff.), Viebahn (2017) and García-Carpintero (2018).

Now, one might think that if indirect assertion is possible in linguistic communication, the same should hold for pictorial communication. Caricatures are perhaps the best candidates here. In caricatures, features of persons or objects are exaggerated – in a way that is comparable to hyperbole in the linguistic case. If caricatures can be used to assert, as does not seem too far-fetched, there is the possibility of c-indirect pictorial assertion.[17] Here is a candidate for an indirect pictorial assertion involving metaphor, analogous to Bergmann's case discussed above: imagine a drawing of a nuclear reactor but with a giant bomb countdown timer wrapped around it. If one thinks that (1) can be used to assert that the nuclear reactor is likely to fail, one might also be sympathetic to the idea that one can assert the same thing with the drawing.

In c-indirect pictorial assertion, the accuracy of a picture and the truth of an assertion made with the help of that picture might come apart: an inaccurate caricature might be used to make a true assertion. After all, the nuclear reactor drawn in the caricature may be likely to fail even though it has no giant timer wrapped around it. Note that this pictorial assertion may well be c-indirect without being t-indirect: the target scene of the caricature may well be the same as the target of the assertion that the reactor is likely to fail.

It may be possible to capture the exaggerations of caricatures without accepting the possibility of c-indirect pictorial assertion. For instance, one could argue that the content put forward is directly expressed by appealing to a certain system of depiction that is operative in caricatures. However, we take it to be more natural to account for cases of this kind in terms of c-indirectness, involving two distinct contents, although we cannot fully defend that view here. Instead, we want to move on to the question of whether there are t-indirect pictorial assertions.

Recall that t-indirect pictorial assertions are assertions where the target differs from the target of the picture that is used to make the assertion. Consider the following example:

> *Outdated Photo*. Sarah is looking after Sue's banana plant while Sue is on a holiday. Sue and Sarah are messaging, and Sue asks Sarah: "Is my banana plant hydrated?" For the last three days, Sarah forgot to water the plant, and so it is not at all well hydrated: its leaves are drooping and showing some brown patches. However, Sarah doesn't want Sue to know about this, as she thinks she can get the plant back into shape before Sue returns. She remembers that five days ago she took a photo of the plant while it was well hydrated, and she sends the photo to Sue.

By sending the photo, Sarah plausibly asserts that the banana plant is well hydrated, and her assertion is false.[18] But the photo Sarah sends is unedited and in an intuitive

[17] Schier (1986, 173) remarks that caricature is "fully analogous to figures of speech - especially hyperbolic speech", though he does not discuss whether caricatures can be used to make pictorial assertions. For an insightful discussion of metaphorical uses of pictures, see Kulvicki (2020, Chap. 6).

[18] In the study conducted by Viebahn and Wiegmann (2023), subjects judged cases of this kind to be lies to the same extent as they judged cases involving edited pictures. In other words, there was no significant difference in lie-ratings between t-direct and t-indirect pictorial assertions. It would be interesting to find out whether subjects judge the *pictures* in cases such as *Outdated Photo* to be accurate or inaccurate.

sense accurate. After all, it is a photo of (the scene of) Sue's banana plant five days ago, and at that time the plant was indeed hydrated. Sarah's assertion, by contrast, is false, since Sarah *presents* the photo as being of (the scene of) the banana plant roughly *at the time she messages the photo*, which it is not. In other words, the target scene of the photo differs from the assertion's target. Thus, the assertion is t-indirect. As a result of this indirectness, the assertion is false even though the picture is accurate. Note that what Sarah asserts is plausibly also expressed by the photo, and so the assertion is c-direct despite being t-indirect. This indicates, again, that the two types of (in)directness may come apart.

Assuming that Sarah's assertion in *Outdated Photo* is indeed t-indirect, it follows that the targets of pictorial assertions are not straightforwardly determined by the targets of the pictures used. This raises the question how the targets of pictorial assertions are determined instead.[19] One view would be that it's the asserter's intentions that determine the target of a pictorial assertion: Sarah intends to make a claim about the state of the plant as it is at the time of messaging the photo, and so the target is the scene of the plant at that time. Though we cannot go into details here, it strikes us as plausible that the asserter's intentions play a crucial role in determining the assertion's target. Importantly for our purposes, such an intentionalist view helps explain how the assertion's target can come apart from the target of the photo used, assuming that the photo's target is already determined when the photo is taken.[20]

How are c-indirectness and t-indirectness related to the traditional (linguistic) notion of an *indirect speech act* (Searle, 1975; Asher & Lascarides, 2001)? Roughly, an indirect speech act is a speech act that is made by making some other speech act, e.g. a request to turn up the heating that is made by asserting that it is chilly. In such cases, the content of the indirect speech act differs from the content of the sentence uttered. It is easy to see that t-indirectness does not suffice to turn a pictorial assertion into a pictorial analogue of an indirect speech act: in cases such as *Outdated Photo*, only one content is involved, and it is also plausible that only one communicative act is performed.[21] By contrast, c-indirect pictorial assertions with caricatures do involve

[19] Thanks to two anonymous referees for raising this question.

[20] More on this assumption in the next section. As two anonymous referees have pointed out, the questions under discussion in our examples (e.g. "Is my banana plant hydrated?") could also be taken to determine a pictorial assertion's target. More would need to be done to spell out how a question determines a target. We want to leave open how this could be done and how plausible the resulting view would be. (See Schwarz (2009, 136–143) for discussion of how a question could determine a topic situation, based on a proposal by Kratzer (see Kratzer, 2019, Sect. 8).) Topic situations arguably don't have a built-in perspective, as Greenberg's target scenes do, and so there's a remaining issue of whether and how questions can determine a target, or topic, that has such a perspective.) Note that even on the intentionalist view, the question under discussion might well play a role in constraining which intentions are reasonable or cooperative, given the aims of the conversation. Important, again, is that if the question under discussion determines the target, that could also explain how a pictorial assertion's target may come apart from the target of the photo used, since there's no guarantee that the question asked is about the same scene as the scene the photo is of. Therefore, we don't need to settle the dispute between an intentionalist and a question-based view of target determination for present purposes.

[21] In *Outdated Photo*, Sarah's t-indirect pictorial assertion is a response to a question, so one might wonder whether t-indirectness in pictorial assertion is related to the linguistic phenomenon of indirectly answering a question. For instance, if John is asked what he is doing, he might indirectly provide the answer that he is studying by saying: "I have an exam on Monday." Again, our sense is that such indirect answers in the lin-

two distinct contents (namely, the content of the picture and the content that is conveyed indirectly, such as, in our example, that the nuclear reactor is likely to fail), and so appear to be closer to typical cases of indirect speech acts. Whether they are exactly analogous to indirect speech acts depends on whether or not one takes them to involve two communicative acts (e.g. an assertion of the exaggerated content and an assertion of the indirectly expressed content). In our view, it is more plausible to hold that only one communicative act is performed in such cases.

5 Repurposing

In the last section, we suggested that Sarah's pictorial assertion in *Outdated Photo* is t-indirect. We take this description to be intuitively plausible. Sarah's assertion seems false even though the photo seems accurate. Our description nicely captures this impression. In this section, we want to argue against a different way of thinking about cases of this kind.

An alternative description of what happens in *Outdated Photo* would be that, by sending the old photo, Sarah *repurposes* the photo and thereby changes the photo's target: when the photo is taken, its target is the scene of the plant at the time of taking the photo, but when the photo is sent five days later, its target has been changed and is now the scene at the time of sending the photo.[22] On this approach, which we will call the *repurposing-approach*, Sarah's assertion is not t-indirect, for the target of the assertion is the same as the photo's (new) target scene. However, we take this alternative approach to be less plausible than our way of accounting for the case.

On the repurposing-approach, the photo Sarah sends to her friend is itself inaccurate with respect to the proposition that the banana plant is well hydrated. As a result, a first cost for this approach is that it seems to have difficulties in accounting for the intuitive difference between *Outdated Photo* and the following example:

> *Edited Photo*. Sarah is looking after Sue's banana plant while Sue is on a holiday. Sue and Sarah are messaging, and Sue asks Sarah: "Is my banana plant hydrated?" For the last three days, Sarah forgot to water the plant, and so it is not at all well hydrated: its leaves are drooping and showing some brown patches. However, Sarah doesn't want Sue to know about this, as she thinks she can get the plant back into shape before Sue returns. She takes a picture of the plant and edits it so that it looks well hydrated. Then she sends the edited photo to Sue.

In this case, Sarah also asserts that the banana plant is well hydrated, but we don't have the same impression that the photo Sarah sends is an accurate photo. The photo itself has been edited and has therefore become inaccurate. But if we take *Outdated*

guistic realm are not analogous to t-indirect pictorial assertions (of the kind at issue here): the case of John and the exam seems to involve two contents, while *Outdated Photo* intuitively involves only one content. Many thanks to an anonymous referee for raising this issue and for suggesting the example.

[22] This notion of *repurposing* is Greenberg's (2018, 873, n. 13), who briefly discusses this possibility.

Photo to be a case of repurposing, then both cases involve inaccurate photos. It does seem to be the case that the two examples are interestingly different and that this difference has to do with the accuracy of the photos used. On the repurposing-approach, however, both assertions involve an *in*accurate photo (that used to be accurate). The only difference is *how* the photos lost their accuracy.

A second cost of the repurposing-approach is that it is incompatible with the view that the content and the target of a picture are fixed by the context of *creation*. Above, we indicated that Greenberg accepts this view: the target scene of a photo, for example, is fixed by the function and setting of the camera at the time the picture is taken.[23] On the repurposing-approach, the context of *use* fixes the target of a picture, at least on some occasions. This result is surprising and at least somewhat counterintuitive. As we described Sarah's communicative act in the second case, she uses a photo of (the scene of) the banana plant *five days ago* to make an assertion about the current state of the plant. But the repurposing-approach entails that this description is mistaken: the target of the photo is (the scene of) Sue's banana plant *at the time of the message*. Relatedly, it seems that no aspect of the photo itself changes when Sarah sends it to Sue, but the repurposing-approach entails that it changes its target scene and that it is transformed from an accurate photo into an inaccurate one.[24] For those reasons, we think that Sarah's assertion is better understood as a t-indirect assertion.[25]

The same diagnosis, we think, applies to the *Eiffel Tower* cases that we sketched in the introduction. Recall that in *Eiffel Tower 2* and *Eiffel Tower 3*, Sybil has an office in Paris but the Eiffel Tower is not visible from her desk. In *Eiffel Tower 2*, Sybil takes a photo sitting at her own desk, edits it carefully so that the Eiffel Tower appears to be visible from her desk, and sends the edited photo to her friend. In *Eiffel Tower 3*, Sybil sends a photo of the view from her colleague's desk to her friend. In both cases Sybil (falsely) asserts that the Eiffel Tower is visible from her own desk. But the two cases are markedly distinct. This distinctness can be captured by our preferred approach that allows for t-indirectness: the assertion in *Eiffel Tower 3* is t-indirect, since the assertion's target (the scene visible from her own desk) differs from the photo's target (the scene visible from her colleague's desk). By contrast, the assertion in *Eiffel Tower 2* is still false, but it is not t-indirect: the photo's target is the same as the assertion's target. So, an account that allows for t-indirectness in pictorial assertion can account for the intuitive difference between the two cases by pointing to a difference in t-directness. An explanation of this kind does not seem to be available on the repurposing-approach.

[23] As mentioned above, similar views have been defended by Lopes (1996), Hopkins (1998) and Abell (2009).

[24] The repurposing-approach seems even less plausible given a two-part model of pictorial accuracy, which requires us to say that the *content* of the photo changes when Sarah repurposes it. Could the content of a photo really change without any editing?

[25] When Sarah sends the photo to Sue in *Outdated Photo*, a new picture token is created on Sue's phone, but we don't think that this constitutes the creation of a new photo (type); the same photo can now be found on two different devices, and this photo was created when Sarah first took it five days ago.

6 Why directness matters

In the previous sections, we introduced the notion of t-indirectness in pictorial assertion and argued against the repurposing-approach as an alternative way of capturing the cases discussed. Now we want to explore some further reasons why t-(in)directness is of interest. Firstly, it highlights the possibility of using pictures not just as a means to express and assert content, but also to provide evidential support for the content asserted. Secondly, it opens up questions about how t-indirectness in pictorial assertion can lead to false beliefs. And finally, it raises the question whether t-indirectness can be found in linguistic assertion.

6.1 Photos as evidence

As long as pictorial assertions are true, why should one care whether they are t-direct or not? That we at least sometimes do care may be illustrated by the following example:

> *Durable Photo.* Sarah is looking after Sue's banana plant while Sue is on a holiday. Sue and Sarah are messaging, and Sue asks Sarah: "Is my banana plant hydrated?" Sarah has diligently watered the plant and so it is well hydrated. The camera on Sarah's phone has stopped working yesterday, but she has a photo of the plant that she took five days ago. Since the plant still looks the same, Sarah sends the old photo to Sue.

Plausibly, Sarah makes a true pictorial assertion. By sending the old photo to Sue, Sarah asserts that the plant is hydrated, which is true. However, we do have the impression that there is something fishy about Sarah's assertion, an impression we would not have if Sarah had sent a current photo instead of an old one. If our diagnosis is correct, the difference is that Sarah's assertion in *Durable Photo* is t-indirect, whereas the assertion would be t-direct if Sarah had sent a current photo. The photo's target in *Durable Photo* is the scene of the plant as it was five days ago, whereas the assertion's target is the scene of the plant as it is at the time of Sarah's message. But why does the t-indirectness strike us as problematic? Why might Sue feel misled if she later found out that the photo was taken days before Sarah sent the message?

We suggest that when it comes to pictorial assertions, we commonly assume that the picture is presented as evidence for what is asserted, as evidence that the assertion is true. This is so in particular when it comes to photos. We generally take photos to provide good evidence for what their target scene is like.[26] Therefore, if someone makes an assertion with a photo, rather than a linguistic assertion, one plausible reason for their choice is that they offer the photo in support of their assertion. But for a picture to provide good evidence for what is asserted, it is not enough that the picture provides good evidence about *its* target: it has to provide good evidence about the *assertion's* target. If the target of the picture is the same as the assertion's target, the

[26] For a discussion of why this might be so, also in comparison with other kinds of pictures, see Abell (2010).

photo does provide good evidence about the assertion's target if it provides good evidence about its own target. But if the target of the picture is not the same as the assertion's target, there is generally no guarantee that the assertion's target and the picture's target are alike in relevant ways. They could differ greatly, in which case the picture would not provide good evidence about the assertion's target even if it did provide good evidence about its own target.[27] This could explain why we often expect that assertions made with photos are not just true but also t-direct: we expect that the photo provides evidence that the assertion is true, and we therefore also expect that the photo's target is the same as the assertion's target. Therefore, if we expect that the photo used to make a pictorial assertion provides evidence that the assertion is true, we usually also expect that the assertion is not merely true but also t-direct.

In cases such as *Durable Photo*, it is plausible that the photo is indeed used not just as a representational device that happens to express the content the asserter wants to express, but is also offered as evidence in support of the asserted content. Sarah could have just as well sent a text saying "Yes, the plant is hydrated", but she decides to send a photo nevertheless. Therefore, we plausibly assume that Sarah offers the photo as evidence, and so we expect the assertion to be not just true but also t-direct. That the photo is accurate with respect to its own target (that lies five days in the past) may be some weak evidence that the plant is still hydrated when Sarah sends the photo, but it is not nearly as strong as a photo taken right before Sarah sends it to Sue.[28]

Sometimes, the circumstances make clear that a picture is not offered as evidence. We expect that t-indirectness matters less to us in such cases. If Sarah had sent a drawing of the plant that she made five days ago, rather than a photo, the impression that there is something fishy going on is much weaker. A drawing usually takes a lot more time and skill to make than it does to take a photo. And a drawing could be made from memory. It's therefore expected that the assertion is t-indirect, and it's clear that the drawing in itself doesn't provide very strong evidence that the plant is currently hydrated. In accordance with that, the assertion doesn't seem problematic despite its t-indirectness.[29]

The difference between pictures that are offered as evidence and pictures not offered as evidence can also help to explain an interesting observation about how pictorial assertions are challenged. As an anonymous referee noted, linguistic asser-

[27] Abell (2010, 89) provides a good illustration of what may happen if one is mistaken about a photo's target scene: "Suppose the photographs apparently showing the first moon landing were in fact taken in a Hollywood studio. While our beliefs about photographic processes may justify an attribution of high EV [epistemic value] to such photographs, it would be wrong to construe them as providing good evidence that a moon landing occurred, since their external object [i.e., roughly, their target scene] is not a moon landing."

[28] This illustrates that how good the evidence is that a photo provides for the truth of a t-indirect assertion depends on a variety of factors that are specific to each case. In this case, a photo that is a month old is worse evidence than a photo that is a week old, which is worse evidence than a photo that is a day old, and so on. Accordingly, we expect that how much t-indirectness we are willing to tolerate also varies from case to case.

[29] Note, however, that certain non-photographic pictures can be used as evidence, too. For example, if Sue were to trace the outlines of the leaves of the banana plant on pieces of paper, or if she were to make carbon copies, then sending the drawings or carbon copies would arguably be a case of asserting and presenting the image as evidence for the assertion.

tions can usually be challenged with replies such as "How do you know that?" (see e.g. Unger, 1979, 263–264 and Williamson, 1996, 505–506), while there are many pictorial assertions for which such a challenge would be unnatural. For instance, in a continuation of *Eiffel Tower 1*, it would be unnatural for the friend to challenge Sybil's pictorial assertion with: "How do you know the Eiffel Tower is visible from your desk?" Attending to the evidential role of photos can help to explain this observation: in sending the photo, Sybil already conveys how she knows that the Eiffel Tower is visible from her desk, and she provides evidence for her knowledge. Nonetheless, there are circumstances in which even photographic pictorial assertions can be challenged in ways typical for assertion. For instance, if the friend has reasons to doubt Sybil's honesty or her reliability (in sending the right picture), it would not be unnatural to challenge Sybil's assertion with: "Are you sure the Eiffel Tower is visible from your desk?" or "Do you really know the Eiffel Tower is visible from your desk?"

In sum, whether a pictorial assertion is t-direct does matter if the picture is offered as evidence in support of the assertion. The reason is that whether a picture supports the assertion depends, among other things, on how the picture's target and the assertion's target are related. In cases where the photo is not plausibly offered as evidence, t-directness seems to matter much less.

6.2 Believing t-indirect assertions

When someone makes a t-indirect pictorial assertion and the addressee believes the asserter, does the addressee always, or at least normally, acquire a *false* belief as a result, due to the indirectness?[30]

Sometimes, as in *Outdated Photo*, a trusting addressee does plausibly acquire a false belief. In that case, if Sue trusted Sarah's t-indirect assertion, she would come to believe that the plant is hydrated, which is false. Other cases are less clear. In *Durable Photo*, Sarah also sends an old photo to Sue, to assert that the plant is hydrated. If Sue believes Sarah, she again comes to believe that the plant is hydrated. But since, in this case, the plant is in fact hydrated, Sarah's assertion is true and so is plausibly also Sue's belief. As discussed in the previous section, Sue has a good reason to assume that Sarah's assertion isn't t-indirect, and so she might also come to believe (falsely) that the photo is current. This kind of false belief may be a common side effect of trusted t-indirect assertions, but it need not always arise, since t-indirect assertions don't always have to pretend to be direct in order to be credible (as the brief discussion of drawings in the previous section suggests).

Note that even if Sarah, in *Durable Photo*, intentionally misleads Sue into thinking that the photo is current, she doesn't seem to be lying. This impression can be explained on the view (mentioned in Sect. 2) that lying requires asserting something one takes to be false. Sarah merely (truly and sincerely) asserts that the plant is hydrated; she doesn't assert that the photo is current, and so she isn't lying about that. (In *Outdated Photo*, by contrast, Sarah is plausibly lying, and she does think that the plant isn't hydrated, contrary to what she asserts.)

[30] Thanks to an anonymous referee for raising this question.

A further question concerns the nature of the beliefs acquired if one trusts a (t-direct or -indirect) pictorial assertion. As an anonymous referee points out, such beliefs could have fine-grained contents and targets (say, sets of viewpoint-centred worlds and viewpoint-centred worlds, respectively), or they could have coarse-grained contents and targets (say, sets of worlds and worlds, respectively) but incorporate the assertion's target into the coarse-grained content. Roughly, on the first view, Sue would come to believe that the plant is hydrated, and she would believe this content of the scene of the plant at the time Sarah sent her message; and on the second, she would come to believe that the plant is hydrated at the scene of the plant at the time Sarah sent her message, and she would believe this content of the actual world.[31]

The former view has the benefit of keeping simple the connection between the contents of pictorial assertions and the contents of the beliefs formed when one believes those assertions (by not requiring that fine-grained contents and targets are mashed into coarse-grained contents and targets). The latter view complicates this connection but keeps the contents of the resulting beliefs simple (by not requiring fine-grained belief contents). More would need to be done to assess whether there are independent reasons for or against the view that when we believe pictorial assertions (or believe pictures, for that matter), our beliefs have fine-grained, viewpoint-sensitive contents. It would go beyond the scope of our paper to settle this question.

6.3 Linguistic indirectness

So far, our focus has been on pictorial assertion and, in particular, on the phenomenon of t-indirect pictorial assertion. As we suggested in Sect. 4.3, c-indirectness in the pictorial realm is akin to non-literalness in the linguistic realm. But can linguistic assertions also be t-indirect?

A linguistic assertion would be t-indirect just in case the assertion's target were to differ from the target of the sentence used to make the assertion. We can adapt the example involving an outdated photo to construct a candidate case for linguistic t-indirectness:

> *Outdated Voice Message.* Sarah is looking after Sue's banana plant while Sue is on a holiday. Sue and Sarah are messaging, and Sue asks Sarah: "Is my banana plant hydrated?" For the last three days, Sarah forgot to water the plant, and so it is not at all well hydrated: its leaves are drooping and showing some brown patches. However, Sarah doesn't want Sue to know about this, as she thinks she can get the plant back into shape before Sue returns. She remembers that five days ago, when the plant was still hydrated, she recorded but forgot to send a voice message in which she says: "Your banana plant is hydrated." So, she goes ahead and sends the old voice message to Sue.[32]

[31] The first kind of view resembles Lewis' (1979) account of *de se* beliefs; according to Lewis, the contents believed are properties (centred worlds) that a believer ascribes to herself. Ross (1997, Chap. 5) points out the parallel between the viewpoint-sensitive contents of pictures and the fine-grained contents Lewis assumes for beliefs.

[32] Thanks to an anonymous referee for suggesting this example.

Sarah makes use of an old message to respond to Sue's question. Sue can't tell that the voice message is five days old, but even so, the old message, one might think, is still true when Sarah sends it to Sue. It was true five days ago that the plant was hydrated, and the message was about the state of the plant five days ago. None of this changes when Sarah sends it to Sue now. And so one might be tempted to think that the old message has itself a certain content and a certain target that are not affected when Sarah decides to send it later. Considering the close structural similarity of this case and *Outdated Photo*, we take Sarah's communicative act to be a good candidate for a t-indirect linguistic assertion.

However, whether Sarah in fact makes a t-indirect assertion by sending the old voice message depends on a variety of additional issues. In the remainder of this section, we will focus on pointing out what we take to be the main factors to consider, rather than on defending a particular verdict.

To begin with, for t-indirect linguistic assertion to be possible, it has to be the case that *sentences* themselves have targets, not just the assertions made using those sentences. In the case of photos, it is very natural to assume that the photo itself has a target, independently of whether it is used by someone to perform communicative acts. In the case of sentences, this way of thinking seems to us to be much less natural. But it may be possible to point to *tokens* of sentences in this respect, so that a particular recording or inscription of a sentence can be said to have a target. The fact that the voice message seems true in the above example seems to support this view, but of course much more would have to be said to spell out and defend such a view of sentence-tokens and their targets.

The example also raises questions about the nature of the targets of sentences (or sentence-tokens) and linguistic assertions. If the target of a sentence is something that doesn't (easily) vary from case to case, t-indirect linguistic assertion can hardly arise. For example, consider the view that these targets are possible worlds. Then the target of Sarah's old message and the target of any assertion she might make by resending the old message are plausibly the same, namely, the actual world. As in the pictorial case, targets in the linguistic realm need to be more fine-grained entities, such as the possible situations of situation semantics or the world-time pairs of views positing temporalist propositions.[33] We're sympathetic to the idea that targets are more fine-grained than possible worlds not just in the pictorial realm but also in the linguistic realm, but this is a complicated issue we can't discuss here.

Assuming that sentence (tokens) and linguistic assertions indeed have fine-grained targets, there would in addition be the question of how their targets are determined; after all, t-indirect linguistic assertion is possible only if the targets of sentence tokens and assertions are determined independently, in ways that allow them to differ. This could happen, for example, if the target of a sentence token were determined by the intentions of the token's producer (in producing the token), and the target of an assertion, made with that token, were determined by the intentions of the asserter (in mak-

[33] Greenberg (2018) mentions topic situations as linguistic analogues of pictorial targets. See e.g. Barwise and Perry (1983) and Kratzer (1989) for discussions of situation semantics; see also Kratzer (2019, Sect. 3) and Lewerentz (2021, 24–26) on the role of topic situations in situation semantics. Austin (1950) is often mentioned as predecessor. Kaplan (1989, 502) accepts a temporalist view of propositions and holds that circumstances of evaluation are world-time pairs.

ing the assertion). In most cases, these two intentions would plausibly be the same, but in cases such as *Outdated Voice Message*, they could come apart. Whether such a view is ultimately correct is, however, an open question.

Finally, even if one agrees that targets in the linguistic realm are more fine-grained than possible worlds and that the targets of sentence tokens and of assertions made with those tokens could in principle come apart, one might still object that it's not clear what kind of speech act Sarah performs by sending the old message. Her speech act is a t-indirect assertion only if it is an assertion in the first place. But it's not obvious that by sending the old message, Sarah has asserted anything at all. Another interpretation would be that she has merely made available an old speech act of hers to Sue: she has widened its audience, rather than making a new assertion.[34] While we're happy to grant that this *could* be what is happening, we think that Sarah *could* also be making a new assertion about the plant's current state. Our impression is that whether Sarah's act is intuitively an assertion depends at least in part on her communicative intentions. Since Sue can't tell that the voice message is old, it seems to her as if Sarah is claiming that the plant is hydrated; and if that's also what Sarah intends to do, then that is a strong (though not conclusive) reason to think that this is indeed what happens. If, on the other hand, Sarah doesn't want to make a new claim about the current state of the plant, but rather wants to mislead Sue in a rather intricate way, by tricking Sue into believing that she is making a (new) assertion, then we think it's not far-fetched that she does in fact not make a new assertion by sending the old message. In the former case, Sarah's communicative act would be a new linguistic assertion, made with an old message, about a new target, and so would be a linguistic t-indirect assertion. In the latter case, by contrast, it would be no (new) assertion and therefore no t-indirect assertion either. But again, more would need to be said to defend this view.

So, despite the structural similarity of *Outdated Photo* and *Outdated Voice Message*, and despite the intuitive t-indirectness of the latter example, there are a number of questions that have to be answered before we can conclude that t-indirect linguistic assertion is indeed possible. We take this to be an important area for further research at the intersection of linguistic and pictorial communication. If it turns out that t-indirectness is indeed possible in linguistic communication, that would speak in favour of a view on which pictorial and linguistic communication have more in common than is often assumed, and it would be an example where close attention to pictorial communication can inform linguistic semantics and pragmatics. However, even a negative answer to the question whether there is linguistic t-indirectness would be an interesting result, as it would point towards a so far unnoticed difference between the mechanisms of linguistic and pictorial assertion. It would mean that pictorial assertion allows for a kind of indirectness that linguistic assertion does not allow for. One explanation for this could be that sentences (or sentence-tokens) simply don't

[34] To make an old speech act available to a new audience may itself be a new communicative act. What is going on may be somewhat similar to the act of retweeting an old message; for an interesting account of retweeting, see Marsili (2021). Though note that in the case of a retweet, it's clear to the readers that they are confronted with an old tweet; Sue, by contrast, doesn't know that the voice message she receives from Sarah is old.

have targets fixed independently of their uses in communicative acts, whereas at least some pictures do.

7 Conclusion

We have argued that we need to distinguish between the accuracy of pictures, on the one hand, and the truth of pictorial assertions, on the other. Since pictorial assertions can be indirect, pictorial accuracy is neither necessary nor sufficient for the truth of pictorial assertions. To that end, we pointed out that there are two different kinds of directness, content-directness and target-directness. And we provided arguments that there are pictorial assertions that are indirect with respect to their targets, that is, pictorial assertions where the target of the assertion differs from the target of the picture used. In such cases, the assertion may seem to be fishy, even though it is true. We suggested that this may be because in pictorial assertion, the picture is often used not just as representational device, but also offered as evidence in support of the assertion, which generally requires that the assertion is direct with respect to its target. Target-directness, though not necessary for truth, is therefore an important characteristic of many pictorial assertions. Finally, we considered whether there can be target-indirectness in the case of linguistic assertion, and tentatively argued that it can.

The discussion of these issues brought out several questions that we could not fully address but that are worthy of further research. Which role do intentions, questions under discussion or other factors play in determining which part of a picture's content is asserted in partial pictorial assertion? Which role do they play in determining the target of a pictorial assertion? What is the nature of the beliefs acquired through pictorial assertion? Do sentences have targets and, if so, of what kind? There is reason to be optimistic that the surge in research on pictorial communicative acts and their relation to linguistic communicative acts will lead to a better understanding of these matters. We hope that our paper can make a small contribution to this process.

Acknowledgements Both authors contributed equally to this paper. For helpful comments and discussion, we would like to thank Eliot Michaelson, Alex Radulescu, Alex Wiegmann, Richard Woodward, the participants of the Fall Language Online Workshop 2020, Miguel Hoeltje's research colloquium at Universität Duisburg-Essen and Gerson Reuter's research colloquium at Justus-Liebig-Universität Gießen, as well as anonymous referees and the editors of *Linguistics and Philosophy*.

Funding Open Access funding enabled and organized by Projekt DEAL. During part of the work on this paper, Lukas Lewerentz was generously supported by a Jacobsen Studentship from the Royal Institute of Philosophy and by the Vice-Chancellors' Fund at the University of Oxford. Emanuel Viebahn would like to acknowledge generous support through a Career Development Award of the Berlin University Alliance, an Anderson Fellowship at the University of Sydney and a fellowship at the Center for Advanced Studies at Ludwig-Maximilians-Universität München (as part of the research group "Persuading under Uncertainty: Challenges and Norms of Science Communication").

Declarations

Competing interests The authors have no relevant financial or non-financial interests to disclose.

References

Abell, C. (2009). Canny resemblance. *The Philosophical Review*, *118*(2), 183–223. https://doi.org/10.1215/00318108-2008-041.

Abell, C. (2010). The epistemic value of photographs. In C. Abell, & K. Bantinaki (Eds.), *Philosophical Perspectives on Depiction* (pp. 81–103). Oxford University Press.

Abusch, D. (2020). Possible-worlds semantics for pictures. In D. Gutzmann, L. Matthewson, C. Meier, H. Rullmann, & T. E. Zimmermann (Eds.), *The Wiley Blackwell Companion to Semantics*. John Wiley & Sons. https://doi.org/10.1002/9781118788516.sem003.

Alston, W. (2000). *Illocutionary Acts and Sentence Meaning*. Cornell University Press.

Asher, N., & Lascarides, A. (2001). Indirect speech acts. S*ynthese*, *128*, 183–228. https://doi.org/10.1023/A:1010340508140.

Austin, J. L. (1950). Truth. *Proceedings of the Aristotelian Society, Supplementary Volumes*, *24*, 111–128. https://doi.org/10.1093/aristoteliansupp/24.1.111.

Bach, K., & Harnish, R. (1979). *Linguistic Communication and Speech Acts*. MIT Press.

Barwise, J., & Perry, J. (1983). *Situations and Attitudes*. MIT Press.

Bergmann, M. (1982). Metaphorical assertions. *The Philosophical Review*, *91*(2), 229–245. https://doi.org/10.2307/2184628.

Blumson, B. (2014). *Resemblance and Representation: An Essay in the Philosophy of Pictures*. Open Book Publishers.

Brandom, R. (1983). Asserting. *Noûs*, 17(4), 637–650. https://doi.org/10.2307/2215086.

Carson, T. L. (2010). *Lying and Deception: Theory and Practice*. Oxford University Press.

Casati, R., & Varzi, A. C. (1999). *Parts and Places: The Structures of Spatial Representation*. MIT Press.

De Leon, C. (2022). Pointing to communicate: The discourse function and semantics of rich demonstration. *Linguistics and Philosophy*. https://doi.org/10.1007/s10988-022-09363-0.

Dixon, D. (2020). Lies in art. *Australasian Journal of Philosophy*, *100*(1), 25–39. https://doi.org/10.1080/00048402.2020.1844772.

Eaton, M. (1980). Truth in pictures. *The Journal of Aesthetics and Art Criticism*, *39*(1), 15–26. https://doi.org/10.2307/429915.

Esipova, M. (2021). On not-at-issueness in pictures. *Glossa: A Journal of General Linguistics*, *6*, 83, 1–20. https://doi.org/10.5334/gjgl.1314.

García-Carpintero, M. (2018). Sneaky assertions. *Philosophical Perspectives*, *32*, 188–218. https://doi.org/10.1111/phpe.12116.

García-Carpintero, M. (2023). Lying versus misleading, with language and pictures: The adverbial account. *Linguistics and Philosophy*. https://doi.org/10.1007/s10988-022-09355-0.

Gariazzo, M. (2019). Assertion and assessment sensitivity. *KRITERION*, *60*(143), 355–376. https://doi.org/10.1590/0100-512X2019n14307mg.

Goodman, N. (1968). *Languages of Art: An Approach to a Theory of Symbols*. Bobbs-Merrill.

Greenberg, G. (2013). Beyond resemblance. *The Philosophical Review*, *122*(2), 215–287. https://doi.org/10.1215/00318108-1963716.

Greenberg, G. (2018). Content and target in pictorial representation. *Ergo*, *5*(33), 865–898. https://doi.org/10.3998/ergo.12405314.0005.033.

Hopkins, R. (1998). *Picture, Image and Experience: A Philosophical Inquiry*. Cambridge University Press.

Hyman, J. (2006). *The Objective Eye. Color, Form, and Reality in the Theory of Art*. The University of Chicago Press.

Kaplan, D. (1989). Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–563). Oxford University Press.
Kjørup, S. (1974). George Inness and the Battle at Hastings, or doing things with pictures. *The Monist*, *58*(2), 216–235. https://doi.org/10.5840/monist197458217.
Korsmeyer, C. (1985). Pictorial assertion. *The Journal of Aesthetics and Art Criticism*, *43*(3), 257–265. https://doi.org/10.2307/430639.
Kratzer, A. (1989). An investigation of the lumps of thought. *Linguistics and Philosophy*, *12*, 607–653. https://doi.org/10.1007/BF00627775.
Kratzer, A. (2019). Situations in natural language semantics. *The Stanford Encyclopedia of Philosophy* (Summer 2019 Edition), E. N. Zalta (Ed.), https://plato.stanford.edu/archives/sum2019/entries/situations-semantics/.
Kulvicki, J. (2020). *Modeling the Meanings of Pictures. Depiction and the Philosophy of Language*. Oxford University Press.
Lewerentz, L. (2021). *Situations, Presuppositions, and Contexts*. DPhil thesis. University of Oxford.
Lewis, D. (1979). Attitudes de dicto and de se. *The Philosophical Review*, *88*(4), 513–543. https://doi.org/10.2307/2184843.
Lopes, D. (1996). *Understanding Pictures*. Oxford University Press.
MacFarlane, J. (2011). What is assertion? In J. Brown, & H. Cappelen (Eds.), *Assertion. New Philosophical Essays* (pp. 79–96). Oxford University Press.
MacFarlane, J. (2014). *Assessment Sensitivity: Relative Truth and its Applications*. Oxford University Press.
Marsili, N. (2021). Retweeting: Its linguistic and epistemic value. *Synthese*, *198*, 10457–10483. https://doi.org/10.1007/s11229-020-02731-y.
Novitz, D. (1975). Picturing. *Journal of Aesthetics and Art Criticism*, *34*(2), 145–155. https://doi.org/10.2307/430071.
Pagin, P., & Marsili, N. (2021). Assertion. *The Stanford Encyclopedia of Philosophy* (Winter 2021 Edition), E. N. Zalta (Ed.), https://plato.stanford.edu/archives/win2021/entries/assertion/.
Peirce, C. S. (1934). Judgment and assertion. In C. Hartshorne, & P. Weiss (Eds.), *The Collected Papers of Charles Sanders Peirce, Volume V* (pp. 385–387). Harvard University Press.
Rescorla, M. (2009). Predication and cartographic representation. *Synthese*, *169*, 175–200. https://doi.org/10.1007/s11229-008-9343-5.
Ross, J. (1997). *The Semantics of Media*. Springer.
Schier, F. (1986). *Deeper into Pictures: An Essay on Pictorial Representation*. Cambridge University Press.
Schlenker, P. (2018). What is super semantics? *Philosophical Perspectives*, *32*, 365–453. https://doi.org/10.1111/phpe.12122.
Schlöder, J., & Altshuler, D. (2023). Super Pragmatics of (linguistic-)pictorial discourse. *Linguistics and Philosophy*. https://doi.org/10.1007/s10988-022-09374-x.
Schwarz, F. (2009). *Two Types of Definites in Natural Language*. Open Access Dissertation. University of Massachusetts Amherst.
Searle, J. R. (1975). Indirect speech acts. In P. Cole, & J. L. Morgan (Eds.), *Syntax and Semantics. Volume 3. Speech Acts* (pp. 59–82). Academic Press.
Stalnaker, R. (1999). *Context and Content. Essays on Intentionality in Speech and Thought*. Oxford University Press.
Stokke, A. (2018). *Lying and Insincerity*. Oxford University Press.
Unger, P. (1979). *Ignorance. A Case for Scepticism*. Oxford University Press.
Viebahn, E. (2017). Non-literal lies. *Erkenntnis*, *82*, 1367–1380. https://doi.org/10.1007/s10670-017-9880-8.
Viebahn, E. (2019). Lying with pictures. *British Journal of Aesthetics*, *59*(3), 243–257. https://doi.org/10.1093/aesthj/ayz008.
Viebahn, E. (2020). Lying with presuppositions. *Noûs*, *54*(3), 731–751. https://doi.org/10.1111/nous.12282.
Viebahn, E. (2021). The lying-misleading distinction: A commitment-based approach. *The Journal of Philosophy*, *118*(6), 289–319. https://doi.org/10.5840/jphil2021118621.
Viebahn, E., & Wiegmann, A. (2023). An empirical perspective on pictorial lies. *OSF Preprint*. https://doi.org/10.31219/osf.io/wp7va.
Walton, K. (1983). Fiction, fiction-making, and styles of fictionality. *Philosophy and Literature*, *7*(1), 78–88. https://doi.org/10.1353/phl.1983.0004.

Williamson, T. (1996). Knowing and asserting. *The Philosophical Review*, *105*(4), 489–523. https://doi.org/10.2307/2998423.
Wolterstorff, N. (1980). *Works and Worlds of Art*. Clarendon Press.