Recycling Spatial Representations for Temporal Interpretations: Korean Verb Compounding Constructions[1]

Vincent Czarnecki — *Rutgers University*
Dorothy Ahn — *Rutgers University*

Abstract.

It is known that there is a strong crosslinguistic tendency to use spatial expressions to explain temporal concepts, as in sentences like *The meeting was pushed ahead two hours* and *This semester flew by so quickly*. This paper offers a novel analysis of a verb compounding construction in Korean, where verbs of motion *o* ('come')/ *ka* ('go') can form compounds with non-motion verbs, yielding a particular temporal interpretation. The analysis uses a modified version of the path-algebraic approach to motion as outlined in Zwarts (2005, 2017). In the analysis, the underlying spatial representations of motion-verbs are recycled to be interpreted temporally, where the only change is the domain of evaluation. Namely, paths are interpreted with respect to a domain of times rather than a domain of spaces.

Keywords: spatial semantics, path-algebra, event semantics, temporal expressions

1. Introduction

It is a well-known fact that languages make all sorts of use of spatial language for non-spatial interpretations. In this paper, we show some systematic ways in which robust and independently motivated spatial representations are recycled and used for temporal interpretation in Korean verb compounding constructions. While this case study serves as a proof of concept in the temporal case, we also show how this can be naturally generalized to other non-spatial domains such as change of state, change of possession, etc. The discussion will focus on existing analyses and our reanalysis of the spatial representations of the motion verbs 'come' and 'go'; however, examples will show how these analyses are applicable in many other contexts as well.

Consider the following English examples below:

(1) a. The man was pushed ahead two feet.
 b. The meeting was pushed ahead two hours.

To understand the use of 'push' in (1a), this requires an understanding that there is a pushing event, where the man initially occupied some source location and subsequently occupies a goal location, particularly two feet ahead of the source. The meaning of (1b) is similar in that the meeting occupies some initial starting time, and after some event it occupies a new starting time, particularly two hours ahead of the initial starting time.

[1]We would like to Michael Glanzberg, Kristen Syrett, Ryan Walter Smith, and attendees of the Meaning Across languages (MAL) Lab at Rutgers for their insight during various stages of this work. We would also like to thank the anonynous reviewers as well as the attendees of SuB 30 at Goethe-Universität Frankfurt for their helpful comments and feedback.

Note that these meanings are essentially the same, but only differ in the domain of evaluation. In the spatial case, we are interpreting the meaning of 'push' as instantiating an event which changes the *physical location* of some entity, whereas in the temporal case, we are interpreting the meaning of 'push' as instantiating an event which changes the *temporal location* of some entity. In this sense, we maintain a single representation and repurpose it depending on the context in which it appears. To see this parallel more explicitly, observe the visualizations below in Figures 1a and 1b. The technical machinery used here will be explained in more detail throughout the paper, but informally, let L_g represent some contextually salient goal location (set of points), let T_g represent some contextually salient goal timeframe (set of times), and let $p(0)$ and $p(1)$ represent the beginning and end of some path p, respectively.

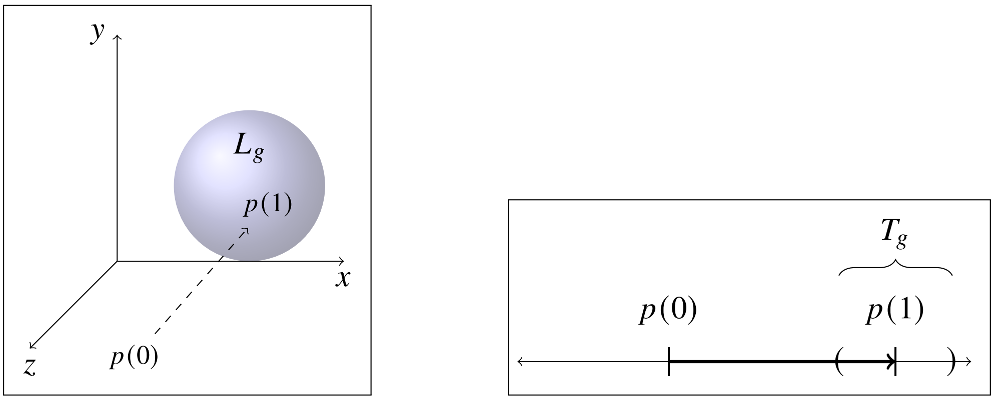

(a) Man-Pushing Event in (1a) (b) Meeting-Pushing Event in (1b)

Figure 1: Pushing Events

We see that these visualizations capture precisely the parallel we wish to encode if Figure 1a is interpreted with respect to a domain of spaces and 1b is interpreted with respect to a domain of times. Below, we preview the corresponding event representations for each example. Informally, the TRACE function returns the (unique) path in a domain of spaces for a given event, where TRACE(e, p) denotes that event e occurs on the path p. This will be covered in more formal detail, but this is sufficient for laying out the parallel here more explicitly.

In (1a), we see that there is a general `motion` event e whose THEME is the referent denoted by the description 'the man', there is a path p ending within some salient goal location L_g (set of points in space), and the spatial distance between the beginning and end of this path is 2 feet.

(2) $\exists e[\text{push}(e) \wedge \texttt{motion}(e) \wedge \textsc{Theme}(e, \text{man}) \wedge \textsc{Trace}(e, p) \wedge p(1) \in L_g$
$\wedge\ |p(0) - p(1)| = 2 \text{ ft.}]$

In (1b), we see that there is a general `motion` event e whose THEME is the referent of 'the meeting', there is a path p ending within some salient goal timeframe T_g (set of points in time), and the temporal distance between the beginning and end of this path is 2 hours.

(3) $\exists e[\text{push}(e) \wedge \texttt{motion}(e) \wedge \textsc{Theme}(e, \text{meeting}) \wedge \textsc{Trace}(e, p) \wedge p(1) \in T_g$
$\wedge\ |p(0) - p(1)| = 2 \text{ hrs}]$

This serves as a simple motivating example displaying the heart of the approach, but it sets the stage very naturally for the primary empirical phenomenon introduced in this paper. In Korean, motion verbs *o* ('come') and *ka* ('go') productively combine with a motion verb to indicate path of the motion (Suh 2000; Choi and Bowerman 1991, a.o.). Consider examples (4) and (5), where *o* ('come') and *ka* ('go') combine with the motion verb 'tule' (*enter*):

(4) Jin-i pang-ey tule-**o**-ass-ta
Jin-NOM room-DAT enter-**come**-PAST-DECL
'Jin came into the room.'

(5) Jin-i pang-ey tule-**ka**-ass-ta
Jin-NOM room-DAT enter-**go**-PAST-DECL
'Jin went into the room.'

These examples are spatial only, where 'tule' (*enter*) instantiates a general `motion` event and the finer perspectival information regarding who is present in the goal location of the motion comes from compounding with *o* ('come') or *ka* ('go'). Interestingly, these Korean motion-path verbs can also combine with *non-motion* verbs such as 'sala' (*live*) and have the spatial interpretation of the path verb temporally reinterpreted. For *o* ('come'), the temporal interpretation is akin to 'do *X* up until time *t*' and for *ka* ('go'), the temporal interpretation is akin to 'do *X* from time *t* on'.[2]

(6) Jin-un kulehkey sala-**o**-ass-ta
Jin-TOP like.that live-**come**-PAST-DECL
'Jin has lived that way.'

(7) Jin-un kulehkey sala-**ka**-ass-ta
Jin-TOP like.that live-**go**-PAST-DECL
'Jin went on living that way.'

Drawing primarily on existing analyses of English 'come'/'go' distinction (Fillmore, 1966; Oshima, 2006; Sudo, 2018) and work on Path Algebras (Zwarts, 2005, 2017), we propose a unified analysis for the Korean examples above. We introduce the notion of a **Spatial Frame**, a body of underlying spatial specifications associated with a given motion verb. Essentially, motion verbs instantiate a general `motion` event, and inherently include a certain amount of spatial information for free, all specified via path restrictions in terms of Path Algebras.

We show that this is not only a useful way to think about motion verbs more generally, but that the Korean cases offer a compelling argument for these sorts of representations given their interpretive flexibility outside of the spatial domain. Namely, work in experimental semantics shows that learners show deep goal-biases across different domains such as change in location, change in state, change in possession, etc. (Lakusta and Landau, 2005). It behooves us then to consider a unified analysis based on a more flexible and abstract notion of paths which can be understood within numerous different structured domains.

[2] We thank an anonymous reviewer for pointing out the resemblance to the English 'went on to *X*', which is similarly interpreted as 'do *X* from time *t* on'. They note a comparative absence of an English counterpart for 'do *X* up until time *t*', but we note that while the interpretation is not exactly the same, it bears similarity in the English 'came to *X*', with an interpretation akin to 'didn't do *X* until time *t*', which we leave for further work.

This paper is organized as follows. Section 2 will provide the formal details of Path Algebras, and give some context for the types of problems they have been employed to solve. This will lead into the notion of a Spatial Frame and the nature of how we view spatial representations in terms of lexical decomposition. We use this to prime the discussion for the distinction between English 'come' and 'go', focusing on existing analyses and laying out our Spatial Frame analysis. Section 3 will give more details regarding the Korean data and our formal analysis. Crucially, this is where we highlight the domain reduction, the central novel contribution of this paper. In Section 4, we show how this style of analysis can be employed to other seemingly disparate empirical phenomena, such as change of state, change of possession, etc. We also briefly remark on some ways of constraining this analysis in terms of logical expressivity, hinting at how this can help delimit the space of logically possible vs. attested lexical meanings for motion verbs. Interestingly while we focus on motion verbs, there is a more general point to be made about the limits of lexicalization in terms of logical expressivity and computational complexity.

2. Spatial frames and motion verbs

In order to understand what kinds of spatial facts are encoded for given lexical items, it is important to consider cases where lexical items have arguments that are interpretable even without showing up explicitly. Consider the examples below from Bhatt and Pancheva (2006) showing examples of more general implicit arguments.

(8) a. Mo noticed __.
b. John won __.
c. John is stronger __.
d. John arrived __.

In these examples, Mo had to have noticed *something*, John had to have won *something*, John has to be stronger *than someone*, and John had to have arrived *somewhere*. In (8d), even without an explicit location present in the utterance, some location is still crucial to correctly interpret the meaning of 'John arrived'. A natural question arises: What is the range of spatial information that a motion verb can lexicalize? We can start by considering the examples below, all of which contain no explicit mention of any location.

(9) Tom left __ although he shouldn't have.

(10) Dick passed __ without being seen.

(11) Harry arrived __ upon realizing he forgot his books.

In (9), 'left' necessitates that the source of Tom's movement is the main location being interpreted. If we fill it with an explicit location, 'Tom left <u>the library</u> although he shouldn't have' can only have an interpretation where his movement started at the library. The intermediate points and endpoints of his movement are irrelevant. In (10), 'passed' necessitates that some intermediate point of Dick's movement is the main location being interpreted. If we have the filled alternative, 'Dick passed <u>the library</u> without being seen', this can only have an interpretation where in the course of his movement, he was located within some proximity of the library. The endpoints of his motion are irrelevant. The use of 'passed' necessitates that some intermediate point of movement is the main location being interpreted. In (11), 'arrived' necessitates that the goal of Harry's movement is the main location being interpreted. If we fill it, 'Harry arrived at the library upon realizing he forgot his books' can only have an interpretation where his movement ended at the library. The beginning and intermediate points of his movement are irrelevant. This shows that reference to particular points, either endpoints or intermediate points along a mover's path, can be lexicalized.

Considering English 'come' and 'go', we see that more than just information about the path can be lexicalized, and that perspectival information regarding the location of particular individuals with respect to the path can also be relevant. Consider the examples below, with the optional overt PP 'to the record store':

(12) a. Steve came __.
b. Steve came to the record store.

(13) a. Steve went __.
b. Steve went to the record store.

Considering (12), there are three scenarios where this could be grammatically uttered. First, this can be true in a context where the speaker was at this location. Second, it can also be true in a context where the addressee was at this location. For example, if I knew that Steve visited you but I had nothing to do with it, I could grammatically utter (12). Finally, this can also be true for a contextually salient third party. Changing the previous context to where I knew Steve would visit my brother, who works at this location, and neither you nor I were present, (12) is still grammatical. Importantly, all three scenarios have in common that Steve travels a path ending at some location (e.g. the record store) and any of the speaker, addressee, or salient third party is located at Steve's point of arrival.[3]

Contrast this with 'go' in (13). Imagining a context where *none* of the speaker, addressee, or salient third party is at the location of Steve's arrival, (13) would instead be grammatical. Suppose on Steve's day off, he takes a trip to this location. If neither I, you, nor some relevant third party are at that location, it would be ungrammatical to utter (12), while (13) is grammatical. However, despite the differences regarding who is present for Steve's arrival, 'come' and 'go' still have a lot in common spatially speaking, the only difference is in whether the perspective holder is present for the mover's arrival.

Thus, we can take the following question as a starting point: how can we lexically encode intermediate/endpoints of a path and the relative location of a given individual with respect to this path? Now that we have given some informal intuition for this question, the discussion will shift to spelling this out formally in terms of Path Algebras and Spatial Frames. The lexical decompositions for these types of motion verbs are given, setting the stage for how these can be interpreted spatially. Then we see that there is only a very minimal jump to be made to yield the non-spatial interpretations of interest to us in this paper. Namely, we see that our construal of paths is much more abstract than simply paths in space, and a simple change of the domain of evaluation gets us exactly what we want.

[3] Some examples later will show that the individual doesn't need to be present *precisely at the time* of the mover's arrival, but to introduce the distinction, this description will suffice.

2.1. Path Algebras and decomposition

In Zwarts 2005, Path Algebras are used to distinguish between different aspectual properties of directional prepositions in English. Informally, directional prepositions can be formalized by specifying restrictions on an infinite set of concatenable paths in a given model of space. The fact that this is an algebraic structure contributes to the primary goal of that paper, which is accounting for telic and atelic directional prepositions regarding cumulativity. This is more or less orthogonal to the discussion here, but the machinery elegantly extends in various ways, which is the focus of this paper. The main elements used for the discussion here are formally defined below.

Following Zwarts (2005), a **path** p is defined as a continuous function from the real unit interval [0,1] to positions in some model of space D_s. More formally, a path is a function $p : [0,1] \rightarrow D_s$. Each point in the interval [0,1] maps directly to a point in the model of space D_s, i.e. a location. It is useful to pick out any point of the path for some intermediate point $0 \leq j \leq 1$, but most crucially, the start of a path p is $p(0)$, and the end is $p(1)$. One could choose a wide variety of domains, but the most straightforward choice is 3-dimensional Euclidean space, or $\mathbb{R}^3$. A **Path Algebra** $\langle \mathbf{P}, + \rangle$ is the set of all possible paths p in the domain of space D_s endowed with a path concatenation operator +, which essentially 'glues' paths together. The details of why this is an algebraic structure are not crucial to the analysis, though informally, it is an algebra because it satisfies a particular set of mathematical properties such as: if you take two paths $p, q \in \mathbf{P}$, their concatenation $p + q$ is still a path in $\mathbf{P}$; paths have *inverses* (the same path backwards); there is an *identity* path (no movement), etc.

So far, we can reference paths and points along them but nothing has been said about how these paths are related to events. There is a sense in which this is useful on its own if we think about modality and stating something about the *possibility* or *necessity* of traversing a path (Martin et al., 2020). Still, this discussion will be limited to cases where path traversal does occur. For this, Zwarts also defines a Trace function, which takes an event e as its argument and returns the (unique) path associated with that event.[4] This Trace function serves to tie events to paths in the domain of space explicitly.[5]

We define a Spatial Frame as the body of underlying spatial specifications inherent to a lexical item. This is a lexical decomposition of the item where each component is some conjunct restricting spatial information. For the cases of dynamic motion under discussion here, these are restrictions on the Trace corresponding to an event. This way of formalizing entails that, even when a verb has no overt locations specified, there is a necessary, underspecified value in its place that come along with it. When they are filled overtly (e.g. 'from the room', 'here', 'there', etc.), this is called 'filling the frame' by inserting lexical information into the underspecified value. The values that can be underspecified include the following: (i) the source, or beginning of the path, (ii) intermediate points on the path, (iii) the goal, or end of the path, and (iv) the location relative to some individual.

[4]This is likened to the spatial trace function as described in Link (1998).

[5]It is worth noting that paths are not explicitly parameterized with respect to time. While this may be expected since it is ontologically necessary that paths are traversed as a function of time, this is irrelevant to how they are *linguistically encoded* since we frequently make use of paths time-independently.

The standard way of representing each of these underspecified values in a lexical item's Spatial Frame is shown in Table 1 below. The term $p(0)\ldots L$ denotes a relationship between the path's beginning and some contextually salient source-location L (sometimes denoted L_s or source-location for convenience). The term $p(j)\ldots L$ denotes that some intermediate point on the path $p(j)$ where $0 < j < 1$ can be in some relation to a particular location L. The term $p(1)\ldots L$ denotes a relationship between the path's end and some contextually salient goal-location L (sometimes denoted L_g or goal-location for convenience). Further, we define a function $Loc(x) : D_e \rightarrow D_s$ which takes an individual $x \in D_e$ and returns their location, an individual point in D_s. Thus, $Loc(x)\ldots L$ denotes that some individual x has some relationship to a particular location L. In this table, L is always assumed to be a subspace of D_s, so a set of points, often determined by the context.

Source	Intermediate Points on Path	Goal	Location Relative to Individual x
$p(0)\ldots L$	$p(j)\ldots L$	$p(1)\ldots L$	$Loc(x)\ldots L$

Table 1: Spatial Frame Specifications

These relationships, '...', can simply be set membership to state that a point $l \in L$, or it can be a more complex relationship, which we will see examples of shortly. At this point, it is also crucial to note that this information is not always asserted and can instead be backgrounded. For example, as we will see in the next section with English 'come', mirroring Sudo (2018)'s analysis, it is backgrounded and not asserted that the perspective holder occupies a shared space with the mover's point of arrival.

For now, restricting the discussion to points on a path, let's revisit examples (9)–(11) to see their corresponding analyses and visualizations. In (9), 'left' necessitates that the source of Tom's movement is the main location being interpreted. If we fill it, 'Tom left the library although he shouldn't have' can only have an interpretation where his movement started at the library. The intermediate points and endpoints of his movement are irrelevant. Thus, there is a `motion` event whose TRACE begins at some relevant source-location L_s, which is analyzed below, along with a corresponding visualization:

(14) $[\![\text{leave}]\!] = \lambda e[\ \ldots\ \texttt{motion}(e) \wedge \textsc{Trace}(e, p) \wedge p(0) \in L_s]$

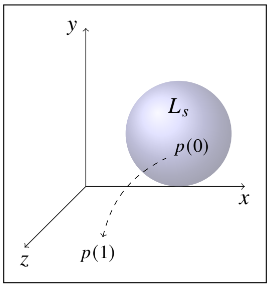

Figure 2: Visualization of $[\![\text{leave}]\!]$

In (10), 'passed' necessitates that some intermediate point of Dick's movement is the main location being interpreted. If we have the filled alternative, 'Dick passed the library without being seen', this can only have an interpretation where in the course of his movement, he was located within some proximity of the library. The endpoints of his motion are irrelevant. The use of 'passed' necessitates that some intermediate point of movement is the main location being interpreted. Here, there is sensitivity to the *proximity* of L as opposed to simply being within L or not. For an intermediate point on the path $p(j)$, one way of saying that $p(j)$ is *near* the location L is if there is an individual point q inside of L such that the distance between $p(j)$ and that point is sufficiently close given some contextually determined threshold.[6] This threshold can be denoted ϵ_c, and so this distance restriction can be stated as $|p(j) - q| \leq \epsilon_c$, meaning that the distance between these two points is less than or equal to the threshold. This analysis actually distinguishes between 'passed *through* __' and 'passed *by* __' since this distance can be zero if the mover is inside the location (if the distance between $p(j)$ and q is zero, they are the same point, meaning that the mover is within L), but it can just as well be any non-zero value, depending on the threshold. This is precisely what the $\leq$ in the distance enforces: if this threshold is $\epsilon_c = 0$, then these two points are the same since $|p(j) - q| = \epsilon_c = 0$, so $p(j) = q$.

(15) $[\![\text{pass}]\!] = \lambda e[\ \ldots\ \texttt{motion}(e) \land \textsc{Trace}(e, p) \land \exists q \in L[|p(j) - q| \leq \epsilon_c]]$

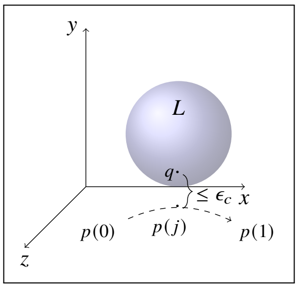

Figure 3: Visualization of [[pass]]

In (11), 'arrived' necessitates that the goal of Harry's movement is the main location being interpreted. If we fill it, 'Harry arrived <u>at the library</u> upon realizing he forgot his books' can only have an interpretation where his movement ended at the library. The beginning and intermediate points of his movement are irrelevant. Thus, there is a `motion` event whose TRACE ends at some contextually relevant goal-location L_g, analyzed as below with a corresponding visualization:

(16) $[\![\text{arrive}]\!] = \lambda e[\ \ldots\ \texttt{motion}(e) \land \textsc{Trace}(e, p) \land p(1) \in L_g]$

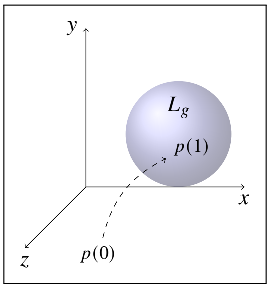

Figure 4: Visualization of [[arrive]]

[6] It is also possible to instead define a NEAR(l_1, l_2) relation that establishes the relative 'nearness' of two locations l_1 and l_2, perhaps evaluated in a similar way, but excluding distance evaluation in the event structure.

Now that our approach to spatial representations has been introduced more formally, seeing the distinction between English ‘come’ and ‘go’ will prime us for a better understanding of how these are utilized in our analysis for the Korean data.

2.2. The semantics of coming and going

Recall from the informal description earlier that the distinction between English ‘come’ and ‘go’ is not necessarily spatial so much as it is perspectival. Namely, English ‘come’ and ‘go’ both have in common that there is movement of a THEME argument from a source location to a contextually relevant goal location. Their difference lies in the relationship to the relevant perspective holder’s location and the mover’s of arrival. Consider the repeated examples below:

(17) a. Steve came __.
b. Steve came to the record store.

(18) a. Steve went __.
b. Steve went to the record store.

What sets apart the ability to grammatically utter (17) and (18) is who is present for Steve’s arrival to the goal location. In (17), this could be uttered if either (i) the speaker spk, (ii) the addressee adr, or (iii) relevant third party *are present*. In (18), this could be uttered if either (i) the speaker, (ii) the addressee, or (iii) relevant third party are *not present*.

Oshima (2006) accounts for this by defining a set *RP*, or Reference Point, which depending on language-specific parameters can include some combination of speaker, addressee and relevant third party. The distinction is in that ‘go’ requires that no member of the RP be at the goal of the described motion, whereas ‘come’ requires that some member of the RP be at the goal. This nicely explains crosslinguistic variation between languages that have tendencies to include or exclude reference to certain individuals using ‘come’ or ‘go’.

Sudo (2018) accounts for the distinction by arguing that ‘come’ has presuppositional content that ‘go’ lacks and derives the implication of ‘go’ as an anti-presupposition. By identifying an individual’s *Home Location* (HL), which is a location in some way canonically associated with a given individual (*born in, lives in, is currently located at, etc.*), the distinction roughly amounts to enforcing that if the goal location is a speaker or hearer’s HL, ‘come’ is licit and otherwise, ‘go’ is licit. The assertive meaning of ‘come’ and ‘go’ are analyzed to be the same, with the only difference being their presupposed content. This subjects them to competition due to MAXIMIZE PRESUPPOSITION! (MP), which favors presuppositionally stronger alternatives. Thus, the choice between ‘come’ and ‘go’ will depend on the speaker or hearer’s HL. More specifically, the speaker will choose ‘come’ if the goal is either the speaker or hearer’s HL, and will choose ‘go’ otherwise. Observe the two examples below with the corresponding analysis from Sudo (2018). Here, $\mathbf{Speaker}^{+}(c)$ refers to the set of individuals including the speaker and anyone associated with them in the context c. Thus, (19) says that if there is some individual x in this set whose HL is London, then the speaker uses ‘come’, and in any other instance they are predicted to use ‘go’.

(19) $((\text{Natasha came to London}))^{g,c_1,c_2} = \lambda k.\exists x \in \mathbf{Speaker}^{+}(c_2)[\mathbf{HL}_{w_{c_2}}(x, \text{London}, t_{c_2})]$

(20) $((\text{Natasha went to London}))^{g,c_1,c_2} = \lambda k.\top$

Here, we see that the path-based analysis detailed above very naturally accommodates these observations. Taking common assertive components of their meaning, there should be a generalized `motion` event whose TRACE ends within some proximity to the location of the perspective holder. In other words, there is a shared contextually relevant goal within which the mover's path ends and the perspective holder are located. To start using our notation, we have an event e for which (i) $\texttt{motion}(e)$, (ii) $\textsc{Trace}(e, p)$, and (iii) $p(1) \in L$. In line with Oshima (2006), depending on the language the perspective holder can be any of the speaker, addressee, or some contextually relevant third party. In our analysis, this individual will be denoted H_p and is assumed to be determined by the context. Our $Loc(x)$ function will be used to reference the location of the perspective holder, $Loc(H_p)$, in relation to a particular point on a motion event's TRACE. With 'come', it will be backgrounded[7] that the persepctive holder is in this space, and all other spatial facts are asserted. In our notation, $Loc(H_p) \in_* L$ states the backgrounded requirement that the perspective holder is occupying L, in proximity to the end of the mover's path since $p(1) \in L$. We assume that 'go' simply lacks this backgrounded information in its lexical decomposition, and similarly to Sudo (2018), MP determines the choice of between the two. Observe the analyses below with the corresponding visualizations, where the only difference is the presence or absence of the backgrounded perspectival information $Loc(H_p) \in_* L$:

(21) $[\![\text{come}]\!] = \lambda e[\ \ldots\ \texttt{motion}(e) \wedge \textsc{Trace}(e, p) \wedge p(1) \in L \wedge Loc(H_p) \in_* L]$

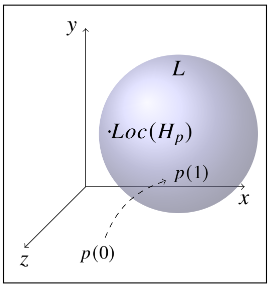

Figure 5: Visualization of $[\![\text{come}]\!]$

(22) $[\![\text{go}]\!] = \lambda e[\ \ldots\ \texttt{motion}(e) \wedge \textsc{Trace}(e, p) \wedge p(1) \in L]$

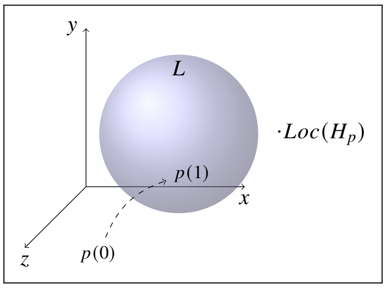

Figure 6: Visualization of $[\![\text{go}]\!]$

Here, $Loc(H_p)$ is explicitly shown outside of L to emphasize the difference between scenarios where one would use 'go' instead of 'come'. One could choose to include $Loc(H_p) \in_* L$ in the definition of 'go', but we follow Sudo 2018 in assuming that this implication arises as an anti-presupposition.

[7]This is notated with $*$ following AnderBois et al. (2015).

We now turn to Korean to show what this path-algebraic view brings to bear in our analysis.

3. Korean verb compounding constructions

In this section, we present a novel analysis of the Korean compounds utilizing the path-based analysis. It is shown how changing the domain of evaluation from a 3-dimensional domain of spaces to a 1-dimensional domain of times accounts for the data in a straightforward manner. In this sense, the independently motivated spatial representations are changed *only in how they are evaluated* for non-spatial interpretations. This not only provides a novel empirical contribution, but also a proof of concept of a much more general method of recycling spatial interpretations.

3.1. Data

Recall that in Korean, motion verbs *o* ('come')/*ka* ('go') can productively combine with a motion verb as in (4) and (5) to indicate path of the motion (Suh 2000; Choi and Bowerman 1991, a.o.).

Also recall that when compounded with *non-motion* verbs, the spatial interpretation of the path verb is temporally reinterpreted. For *o* ('come'), the temporal interpretation is akin to 'do *X* up until time *t*' and for *ka* ('go'), the temporal interpretation is akin to 'do *X* from time *t* on', as in the examples below:

(23) Jin-un kulehkey sala-**o**-ass-ta
Jin-TOP like.that live-**come**-PAST-DECL
'Jin has lived that way.'

(24) Jin-un kulehkey sala-**ka**-ass-ta
Jin-TOP like.that live-**go**-PAST-DECL
'Jin went on living that way.'

(25) Jin-un hankwuk-umsik-ul meke-**o**-ass-ta
Jin-TOP Korean-food-ACC eat-**come**-PAST-DECL
'Jin has eaten Korean food.'

(26) Jin-un hankwuk-umsik-ul meke-**ka**-ss-ta
Jin-TOP Korean-food-ACC eat-**go**-PAST-DECL
'Jin went on eating Korean food.'

This construction is compatible with non-motion predicates across various Vendlerian classes, including statives ('love-come'), activities ('run-come'), achievements ('realize-come'), and accomplishments ('build-come'). With activity predicates, the construction yields a habitual or iterative imperfective interpretation. With achievement and accomplishment predicates, the resulting interpretation necessarily involves iteration over the relevant temporal interval, giving rise to a plurality of achievement or accomplishment events. There is no grammatical restriction on the kinds of predicates that may appear in the construction. The only limitation seems to be whether the predicate can be given an imperfective, habitual, or plural interpretation over a temporal span.

Up until this point, these spatial facts were all understood with respect to a 3-dimensional domain of spaces D_s, which we considered for convenience to be 3-dimensional Euclidean space $\mathbb{R}^3$. Thus, in a Path Algebra, paths p were defined as parameterized curves within this 3-dimensional space. Thinking more generally, there is nothing that precludes us from considering paths in other structured domains. In line with Krifka (1998), we could shrink our domain D_s to a 1-dimensional real line $\mathbb{R}$ so that we can reason about temporal intervals. Paths are still parameterized curves within this space and nothing changes about how we define them, but only how we interpret them.

3.2. Domain reduction

Flattening the 3-dimensional domain of spaces to a 1-dimensional domain of times results in the following interpretive changes, shown in Table 2. In the spatial case, it was relevant to consider a contextually salient goal location L, which is a set of points in the space. From there, we could reason about the location of the perspective holder with respect to the mover's path and this location L. In the temporal case, suppose we have a goal temporal interval T, which is a set of times on a line. From there, we can reason about a contextually relevant time, say the utterance time, t_U. There is a parallel to be drawn here where we can now reason about the runtime of events taking place on this line, and where they are *temporally located* with respect to one another.

Spatial Representation	**Temporal Representation**
salient *space* L (a set of *points*)	salient *interval* T (set of *times*)
salient *location* of perspective holder $Loc(H_p)$	salient *time* (say, utterance time t_U)

Table 2: Interpretive Changes from Domain Reduction

Recall the intended interpretations for (25) and (26), stated informally below:

1. ⟦eat-Korean-food-**come**⟧(x): up until time t, individual x eats Korean food

2. ⟦eat-Korean-food-**go**⟧(x): from time t on, individual x eats Korean food

If t_U is the utterance time, T is a relevant temporal interval, and p is a general temporal path where the runtime of multiple, non-overlapping eating events are sub-paths of p, this corresponds to:

1. ⟦eat-Korean-food-**come**⟧(x): $p(1) \in T$

2. ⟦eat-Korean-food-**go**⟧(x): $p(0) \in T$

Thus, there is a general temporal path p where there is a set of non-overlapping eating events $e_1, \ldots, e_n$ occurring along that path. This means that the runtime of each event, $\tau(e_i)$ for $i \in [1, \ldots, n]$, is a sub-path of this general trajectory path p, denoted $\tau(e_i) \leq p$. Then crucially, (i) if those *end within T* we predict 'eat-Korean-food-**come**' and (ii) if those *begin within T* we predict 'eat-Korean-food-**go**'. More explicitly, there are some distinct events $e_1, e_2, \ldots, e_n$ such that the following are true:

- They are eating events:
  $\forall e_i[eat(e_i)]$
- Their AGENT's are all Jin:
  $\forall e_i[\text{AGENT}(e_i, \text{Jin})]$
- Their THEME's are all some type of Korean food:
  $\forall e_i[\text{THEME}(e_i, x)] \mid x \in [\![\text{Korean-food}]\!]$
- Each event runtime is a *sub-path* of a main trajectory path p:
  $\forall e_i[\tau(e_i) \leq p]$
- They have no overlapping runtimes
  $\forall_{e_i,e_j}[\neg \circ (\tau(e_i), \tau(e_j))]$
- The runtime of the {first/last} event coincides with the {beginning/end} of main trajectory path:
  -come: $\tau(e_n)(1) = p(1)$
  -go: $\tau(e_n)(0) = p(0)$

These are both shown visually below in Figures 7 and 8, respectively:[8]

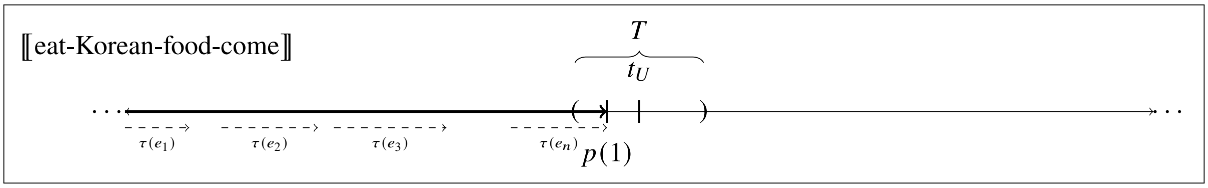

Figure 7: Analysis of (25)

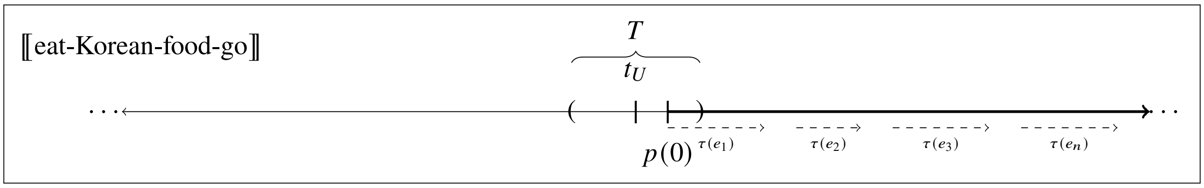

Figure 8: Analysis of (26)

Note the parallel where in the spatial case, we had a path ending in a contextually relevant set of points with some relationship to the location of an individual, and here we have a path ending in a contextually relevant set of times with some relationship to the time corresponding to the utterance. In this sense, we are directly using the spatial meanings of 'come' and 'go' to access the temporal interpretations of these compounding constructions.

There is a crucial point to note here regarding the reference of the beginning and end of this trajectory path. It may appear at first glance that this collapses the meaning of 'go' into a source-oriented predicate rather than a goal oriented predicate, since the sensitivity is with respect to

[8] This bears some resemblance to approaches taken in Deo (2009); Deo et al. (2013) to imperfectives and change of state predicates in terms of non-overlapping sub-intervals and value difference in ordered domains.

the placement of $p(0)$, where in the spatial case it was $p(1)$ for both 'come' and 'go'. This distinction is an illusory one, since the truly important point is about *(non-)inclusion*. More specifically, the spatial case of 'go' was sensitive to non-inclusion of a particular individual in a set of points and in the temporal case of 'go' is sensitive to non-inclusion of particular time in a set of times. While $p(0) \in T$ captures the core of the intended meaning, the deep similarity between the spatial and temporal case really rests on the fact that $p(1) \notin T$ for compounds with 'go'. Another apparent difference is the fact that when discussing locations, it is straightforward to talk about the inclusion or non-inclusion of an individual within a location at a given time, whereas time simply elapses and there is a relative uncertainty about where the temporal path ends up in the future. However, this is a clear case of a spatial meaning being temporally reinterpreted in a non-spatial manner. This is one example, but some natural questions arise: what are the range of ways that spatial representations can receive non-spatial interpretations? Are they always temporal? What are the limits of what is lexically encoded?

4. Further implications

The Korean data offers an empirically and theoretically interesting case of cross-domain reasoning between spatial and non-spatial interpretations. While this is interesting, it poses a deeper question of how far we can take this approach empirically. What other sorts of phenomena can we imagine covering with this method? Lakusta and Landau (2005) notes that learners appear to show a goal-bias in spatial language, but also in the domains of change of state and change of possession.

4.1. Reduction to other non-spatial domains

Consider the example below, where use of 'went' is not spatial but indicates a change in state. Initially, John began at some emotion, happiness, and as a result of some event he ended at another emotion, sadness.

(27) John went from happy to sad very quickly.

In all other cases, we evaluated with respect to domains of space which were real-valued. This is not a strict necessity, we could have evaluated with respect to structured domains of another kind. Suppose we have a partially ordered space of emotions. Defining a path in this structured domain would not be a continuous, parameterized curve in space, but would instead be defined over members in this partially ordered space, emotions in this case.

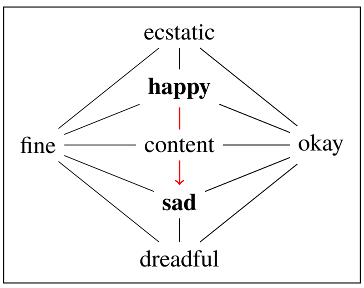

Figure 9: Change of State

From a high level, this retains the more abstract notion of a path, that it begins somewhere in a space, and ends somewhere (else) in that space. We have a path p where $p(0)$ = **happy** represents John's initial state, and $p(1)$ = **sad** represents his final state. This is another case in which the spatial meaning of 'went' is interpreted non-spatially due to a change in the domain of evaluation, but is still defined in terms of its spatial meaning.

Consider the example below, where the use of 'went' is again non-spatial, but refers instead to a change in possession. If we have a partially ordered space of individuals where each member is an individual or collection thereof, paths in this space correspond to stages of ownership. Initially, the tickets belonged to John and then as a result of some event they belonged to Dean and Gene. Note that the arguments can be explicit or implicit and still receive similar analyses as in examples like *The tickets came from Gene* or *The tickets went to John and Gene*.

(28) The tickets went from John to Gene and Dean.

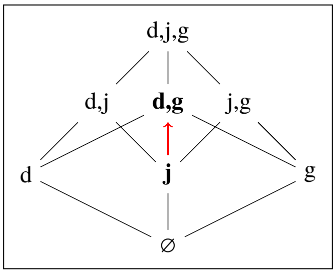

Figure 10: Change of Possession

Coming back to Korean, the example below shows that compounding with a non-motion verb like 'talma' (*resemble*) yields a meaning where resemblance between two individuals grows as time goes on. Namely, initially Jin and Jimin's degree of resemblance could be minimal but overtime it grows more and more.

(29) Jin-kwa Jimin-un talma-**ka**-ss-ta
Jin-CONJ Jimin-TOP resemble-**go**-PAST-DECL
'Jin and Jimin went on resembling each other more.

The nature of what such a domain of space would look like is perhaps less clear than those of the example above, but this is another intriguing case of a different non-spatial interpretation of a compound with a simple motion verb.

4.2. Lexicalization constraints

Taking this approach to spatial representations and reinterpretations thereof, there is a natural question regarding how this system can be constrained. There are many logically possible spatial generalizations one could imagine lexicalizing which appear to be systematically absent from a typological perspective. This is important not only for understanding the nature of spatial representation, but also the limits of how it can be recycled crosslinguistically.

Recent work in computational phonology has highlighted the deep relationship between computational simplicity and typology (Heinz, 2018; Lambert et al., 2021). Namely, different levels of logical complexity yield different *possible vs. attested* phonological generalizations, which is incredibly revealing from a learnability and typological perspective. One common way of displaying this is delineating between string languages defined using First-Order Logic (FO) vs. Monadic Second Order Logic (MSO). Consider a possible phonological well-formedness condition that states a word is well formed iff it has an even number of nasals. This formal language is described below:

(30) EVEN-NASAL: A word is only well-formed iff it has an *even number of nasal consonants*

Ungrammatical:	**Grammatical**:
*pa**n**a	pa**n**a**n**
*pa**m**	pa**n**a**n**a
*pa**n**a**n**a**m**	pa**n**a**n**a**m**a**m**
⋮	⋮

Figure 11: Grammaticality Conditions of EVEN-NASAL

Such generalizations require the power of MSO, and appear to be systematically absent from human language phonology, taken to support the idea that phonological generalizations must be less complex than those generable by MSO.[9] Crucially, this generalization is not inherently about nasality, but about *parity*. Thus, banning levels of expressivity that give rise to parity-based generalizations is not exclusive to phonology, but we can use this as a diagnostic in other linguistic domains as well. Using this as a direct parallel, consider a hypothetical motion verb below in (31) with small visualizations showing well-formed and ill-formed scenarios with respect to the (now informal) definition:

(31) $[\![\text{go}_{even}]\!] = \lambda x \lambda y[x \text{ went to } y \textit{ in an even number of trips}]$

Ungrammatical: **Grammatical**:

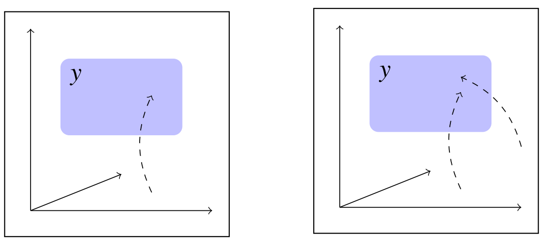

Figure 12: Example Grammaticality Conditions of $[\![\text{go}_{even}]\!]$

Motion verbs sensitive to parity in this way also appear to be systematically absent. This parallel highlights a deep computational similarity across domains, which warrants further investigation and explanation. Leaving a fuller treatment to future work, we give a sketch for this parallel and its potential implications. In FO, one can quantify over *individual elements* x, y, z that satisfy some property φ, whereas in MSO, one can quantify over *arbitrary subsets of elements* $X = \{x_1, \ldots, x_n\}$ that satisfy some property φ. The intuition for EVEN-NAS is that there exists some placeholder set X, where the word's *first nasal* is in X, the word's *last nasal* is not in

[9]This is directly related to what has been referred to as *The Subregular Hypothesis*.

X, and *each nasal alternates* being in X/not in X. Nasality happens to be the sensitivity for this example, but it can be any more general property φ, shown in (32). Using our path-based analysis, φ could just as well express that some corresponding path p ends in the location y. Generally speaking, $\texttt{first-}\varphi(x)$ denotes that an object is the first relevant x such that $\varphi(x)$ in the relevant context, $\texttt{last-}\varphi(x)$ denotes that an object is the last relevant x such that $\varphi(x)$ in the relevant context, and $\triangleleft_\varphi(x, y)$ is a relativized precedence between different objects x, y such that $\varphi(x)$ and $\varphi(y)$. In the case of Even-Nas this would be relativized string precedence (like creating a tier of nasal segments), whereas in the case of $[\![\text{go}_{even}]\!]$ this would be some sort of relativized temporal precedence or containment relation (situating runtimes of distinct events).

(32)
$$\textsc{Even-}\varphi := \exists X[(\forall x)[\texttt{first-}\varphi(x) \rightarrow X(x)] \\ \land (\forall x)[\texttt{last-}\varphi(x) \rightarrow \neg X(x)]] \\ \land (\forall x, y)[x \triangleleft_\varphi y \land (X(x) \leftrightarrow \neg X(y))]$$

With φ standing in for a predicate where a mover's path ends in the location y (e.g. $p_n(1) \in y$ where n represents the n-th trip taken by the mover in the n-th movement event), this definition will crucially only ever be satisfied for a situation wherein a mover ended their path an even number of times. Below, a scenario where it would be false is shown, since it all three conditions cannot be simultaneously true with five paths, as all paths cannot alternate between being in and not in X if $p_1 \in X$ and $p_5 \notin X$:

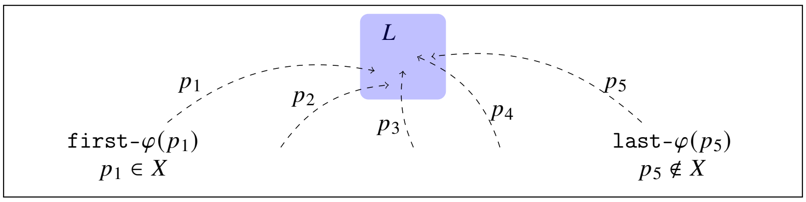

Figure 13: Scenario where $[\![\text{go}_{even}]\!]$ is False

This ability to parity-count necessitates quantification of sets, which might suggest that we eliminate such generalizations via banning quantification of sets, as in the phonological case. However, while this does prune out these unwanted lexicalization patterns, more work remains to be done to determine whether this would accurately reflect the semantic typology.[10] There is a question of whether or not phenomena like pluractionality should require some form of quantification of sets of events. It is possible that a delicate balance could be struck by considering some less complex fragments of MSO, which is left for future work.[11]

These sorts of generalizations also appear to be absent outside of the spatial domain, such as a hypothetical adjectival modifier which means 'a number divisible by 3 of . . . '. However, this serves to show that there are many interesting unresolved questions about the relationship between logical expressivity and typological variation.

[10]There are many other unattested MSO-definable phonological generalizations, such as so-called "First-Last Harmony", where the first and last segments of a word must share a particular quality (e.g. voicing, nasality, etc.). One could imagine a a spatial-event parallel where a lexical item encodes the following: only the first and last paths traveled end at a particular, common point ($p_1(1) = g \land p_n(1) = g$) but the endpoints of all paths temporally between them don't matter ($p_i(1)$ need not be g where $1 < i < n$), which seems like a unnatural thing to lexicalize.

[11]We note that the definitions of both patterns rely on negation, which could be carefully reeled in to adjust complexity and by extension, available generalizations.

5. Conclusion

In this paper, we provided a novel analysis for a Korean verb-compounding construction, wherein compounds with motion verbs receive a temporal interpretation. This analysis was stated in terms of what we call Spatial Frames, using Path Algebras to formally describe the inherent spatial specifications of a given lexical item. The core of our analysis rests on changing the domain of evaluation from a domain of spaces to a domain of times. We further showed how a simple change of the domain of evaluation to other structured domains neatly extends this analysis to other non-spatial interpretations of spatial terms, like change of state, possession, or degrees of a property between individuals. We also identified a typological gap regarding possible vs. attested lexical representations and expressed some preliminary thoughts on how to account for such gaps. Further work could investigate how this style of analysis could be applied to other non-spatial phenomena expressed via spatial language.

References

AnderBois, S., A. Brasoveanu, and R. Henderson (2015). At-issue proposals and appositive impositions in discourse. *Journal of Semantics 32*(1), 93–138.

Bhatt, R. and R. Pancheva (2006). Implicit arguments. *The Blackwell companion to syntax*, 558–588.

Choi, S. and M. Bowerman (1991). Learning to express motion events in English and Korean: The influence of language-specific lexicalization patterns. *Cognition 41*(1-3), 83–121.

Deo, A. (2009). Unifying the imperfective and the progressive: Partitions as quantificational domains. *Linguistics and Philosophy 32*(5), 475–521.

Deo, A. S., I. Francez, and A. Koontz-Garboden (2013). From change to value difference. In *Semantics and Linguistic Theory*, pp. 97–115.

Fillmore (1966). Deictic categories in the semantics of 'come'. *Foundations of Language. International Journal of Language and Philosophy 2*(3).

Heinz, J. (2018). The computational nature of phonological generalizations. *Phonological Typology, Phonetics and Phonology 23*, 126–195.

Krifka, M. (1998). The origins of telicity. In *Events and grammar*, pp. 197–235. Springer.

Lakusta, L. and B. Landau (2005). Starting at the end: The importance of goals in spatial language. *Cognition 96*(1), 1–33.

Lambert, D., J. Rawski, and J. Heinz (2021). Typology emerges from simplicity in representations and learning. *Journal of Language Modelling 9*.

Link, G. (1998). Algebraic semantics in language and philosophy.

Martin, F., M. Grant, C. Piñón, and F. Schäfer (2020). A new case of low modality: Goal pps. In *Semantics and Linguistic Theory*, pp. 562–582.

Oshima, D. Y. (2006). Motion deixis, indexicality, and presupposition. In *Semantics and linguistic theory*, pp. 172–189.

Sudo, Y. (2018). Come vs. go and perspectival shift. *New Brunswick, NJ: Handout for talk at Rutgers University Department of Linguistics*.

Suh, Y. (2000). *A study of the auxiliary verb construction and verb serialization in Korean*. University of Washington.

Zwarts, J. (2005). *Prepositional aspect and the algebra of paths*.

Zwarts, J. (2017). *Spatial semantics: Modeling the meaning of prepositions*.