Feature Geometry and the Person Case Constraint: An Algebraic Link

Thomas Graf
Stony Brook University

ABSTRACT

I give an algebraic characterization of the four attested Person Case Constraints that completely abstracts away from the details of their syntactic encoding. All four constraints are reduced to one general prominence principle, with the cross-linguistic variation stemming from differences in how the individual person features are ranked with respect to prominence. These rankings, in turn, form a natural class in the sense that there is an algebraically natural class of functions that derives them from Zwicky's (1977) person hierarchy. The algebraic perspective thus highlights the feature regularities that underly a specific area of morphosyntax without relying on a specific syntactic analysis. This should make it possible in the future to compare areas of morphosyntax whose analysis builds on vastly different pieces of technical machinery.

Introduction

A lot of work in recent years has focused on morphosyntactic phenomena, in particular the Person Case Constraint (PCC; Kayne 1975, Bonet 1991, Bonet 1994, Anagnostopoulou 2005, Adger & Harbour 2007, Nevins 2007, Béjar & Rezac 2009, Walkow 2012, among others). The PCC is a restriction that renders the grammaticality of direct object (DO) and indirect object (IO) clitic combinations contingent on the person specification of these clitics. The existence of such a constraint has prompted various questions about the feature checking operations in Minimalist syntax, as has the fact that there are several attested variants of the PCC, but fewer than one might expect.

So far, all the work on the PCC has sought to accomplish at least two things:

- provide a syntactic mechanism that is capable of enforcing the PCC, and
- explain the possibility of cross-linguistic variation as well as its limited range.

Intricate technical machinery needs to be invoked in order to derive PCC effects in the syntax — machinery that relies on very specific assumptions about how feature checking takes place and what the specific features and their values look like. This technical accuracy is indispensable for the first problem, the implementation of the PCC, but it also means that the proposed solutions to the second problem are stated in very technical terms that can easily obfuscate important generalizations. This paper is meant to state these generalizations about cross-linguistic variation in as lucid a fashion as possible by deliberately abstracting away from the syntactic encoding of the PCC.

The core insight is that an algebraic perspective reveals the hidden structure of the dependencies that previous analyses of the PCC enforce through various syntactic means. More precisely, one can posit a generalized PCC that bans combinations of clitics based on the locality of their respective features. This locality is encoded in terms of a specific algebraic structure called a *presemilattice*, and cross-linguistic variation arises from language-specific differences in the presemilattice. These presemilattices furthermore have the property that their structure closely resembles the algebra underlying the person hierarchy proposed by Zwicky (1977). Assuming Zwicky's person hierarchy as a universal starting point, then, the range of attested PCCs is characterized in terms of the generalized PCC and the class of locality hierarchies that can be obtained from the person hierarchy in a specific algebraic way.

There are two ways to interpret the algebraic result. The weaker interpretation treats it as merely a useful notation that brings to light certain generalizations that are already inherent in the previous syntactic approaches but are easily missed there among the technical details. The stronger interpretation views it as a genuine claim about language: morphosyntactic constraints operate over scales and hierarchies that correspond to specific algebraic structures, and these structures can be derived in a specific fashion from more primitive objects such as the person hierarchy or some universally fixed feature geometry (cf. Harley & Ritter 2002a,b and Nevins 2007). The stronger stance is more interesting as it predicts that the algebraic dependencies identified here should also play a role in other areas of morphosyntax such as omnivorous number (Nevins 2011), resolved gender agreement (Corbett 1991), and comparative suppletion (Bobaljik 2012). At this point it is unclear whether this prediction is borne out, but considering that current analyses of these phenomena rely on very different technical machinery, an approach that factors out these details is more likely to discover commonalities. So even if one thinks that the algebraic perspective is but a conceptual aid at this point, it holds the promise of more profound generalizations across multiple domains in the near future.

The paper is laid out as follows. Section 1 gives a quick introduction to the PCC and how it differs across languages. After that I present my algebraic characterization via presemilattices in Sec. 2. I then show in Sec. 3 how these presemilattices can be derived in a natural fashion from a single universal structure furnished by Zwicky's (1977) person hierarchy. Section 4 concludes with some additional observations about the predicted typology.

1 What is the Person Case Constraint

The PCC is attested in a variety of languages such as Catalan, French, Spanish, and Classical Arabic, where the grammaticality of direct object (DO) and indirect object (IO) clitic combinations is contingent on the person specification of said clitics (I abbreviate the person features by 1, 2, and 3, respectively). This is illustrated below for French, where a 3IO clitic may combine with a 3DO clitic, but not a 1DO clitic.

(1) Roger **me*/*le* *leur* a presésenté.
Roger 1SG/3SG.ACC 3PL.DAT has shown
'Roger has shown me/him to them.'

Languages differ with respect to the clitic combinations they allow, giving rise to four attested variants of the PCC:

- *Strong PCC* (S-PCC): DO must be 3. (Bonet 1994)
- *Ultrastrong PCC* (U-PCC): DO is less local than IO, where 3 is less local than 2 and 2 less local than 1. (Nevins 2007)
- *Weak PCC* (W-PCC): 3IO combines only with 3DO. (Bonet 1994)
- *Me first PCC* (M-PCC): If IO is 2 or 3, then DO is not 1. (Nevins 2007)

The patterns generated by these constraints are listed in Tab. 1, following the presentation in Walkow (2012). Note that I omit the diagonal because combinations of clitics with identical person specification commonly show special morphological behavior such as spurious *se* in Spanish.

IO↓/DO→	1	2	3
1	NA	*	✓
2	*	NA	✓
3	*	*	NA

(a) S-PCC

IO↓/DO→	1	2	3
1	NA	✓	✓
2	*	NA	✓
3	*	*	NA

(b) U-PCC

IO↓/DO→	1	2	3
1	NA	✓	✓
2	✓	NA	✓
3	*	*	NA

(c) W-PCC

IO↓/DO→	1	2	3
1	NA	✓	✓
2	*	NA	✓
3	*	✓	NA

(d) M-PCC

Table 1: Attested variants of the PCC

The tables clearly reveal the typological puzzle posed by the PCC. Even if one ignores the diagonal, there are six cells that can take one of two values — grammatical or ungrammatical. Without further restrictions there should thus be $2^6 = 64$ variants of the PCC. Yet we only find four. What is it that makes these four variants so special?

2 Algebraic Characterization

2.1 The Generalized PCC and Locality Hierarchies

The definition of the U-PCC above differs from the others in that it consists of two parts: a general constraint "DO is less local than IO" and a metric for computing locality. Without any further assumptions about what counts as a locality metric, this modular way of stating PCCs can trivially be extended to the other three types, as each metric need merely reproduce the relevant relation encoded in Tab. 1. Surprisingly, though, the relevant metrics turn out to be anything but arbitrary.

Slightly rephrasing the U-PCC, we can take the universal component of all PCCs to be given by the *Generalized PCC* (G-PCC): IO is not less local than DO. Denoting "x is more local than y" by $x < y$, the G-PCC can be written IO $\nless$ DO. The crucial parameter is the relation denoted by $<$. It turns out that for each PCC we may equate $<$ with reachability in one of the directed graphs in Fig. 1 such that $x < y$ iff x is reachable from y in G. In the case of the S-PCC, for instance, the G-PCC produces the intended results iff $1 < 2$, $2 < 1$, $3 < 1$, and $3 < 1$. The S-PCC graph is such that 1 is reachable from 2, 2 is reachable from 1, and 3 is reachable from both 1 and 2. Given our definition of < in terms of graph reachability, this gives us exactly the desired locality hierarchy <.

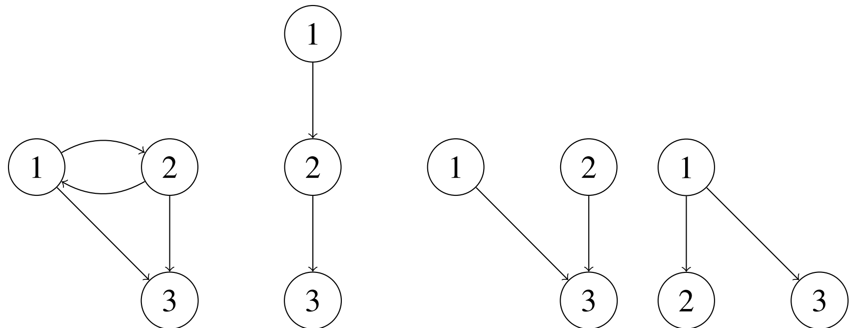

(a) S-PCC (b) U-PCC (c) W-PCC (d) M-PCC

Figure 1: Graphs for the four variants of the PCC

2.2 A First Restriction on Rankings: Preorders

Graph reachability is a relation over nodes, and as such we can classify it according to which relational properties it satisfies. Here are three of the most important ones:

transitivity A relation *R* is transitive iff it holds for all *x*, *y* and *z* that if *x* is related to *y*, and *y* is related to *z*, then *x* is related to *z*.

antisymmetry A relation *R* is antisymmetric iff it holds for all *x* and *y* that if *x* is related to *y*, and *y* is related to *x*, then *x* and *y* are the same node.

reflexvity A relation *R* is reflexive iff every node is related to itself.

For our purposes, transitivity means that if a node *z* can be reached from a node *x* by following one arrow and then another one, there is an arrow directly from *x* to *z*. Antisymmetry states that our graph contains no non-trivial loops — we cannot move from *x* to a different node *y* and then make our way back by following the arrows in the graph. Finally, reflexivity requires every node in the graph to have a trivial loop, i.e. an arrow that points at the node itself.

The graphs in Fig. 1 seem to fail all three properties. In the U-PCC, there is an arrow from 1 to 2 and one from 2 to 3, but no arrow from 1 to 3, so transitivity does not hold, apparently. In the S-PCC, we can follow the arrows from 1 to 2 and back to 2, which violates antisymmetry. And reflexivity does not hold because no node has an arrow pointing at itself. This assessment, however, interprets the graphs in too literal a fashion. For graphs, it is standard to only draw in those parts of the reachability relation that are absolutely necessary. The reachability relation is by default assumed to be transitive, so we do not need to add these arrows to the graphs. Hence transitivity is actually satisfied, it's just not explicitly marked in the graphs. Similarly, reachability is also taken to be reflexive even though our depictions of the graphs do not explicitly show the loops. Overall, then, reachability is transitive and reflexive, but not antisymmetric, and our graphs do in fact omit certain arrows that can be easily inferred from transitivity and reflexivity.

Note that transitivity is an essential property for the graph-based representation of the U-PCC. The U-PCC is the output of the G-PCC operating over the hierarchy 3 < 2 < 1, or more explicitly, the locality rankings 3 < 2, 3 < 1,

$2 < 1$. The graph of the U-PCC, however, only encodes $3 < 2$ and $2 < 1$, the third statement $3 < 1$ is inferred via transitivity of reachability. Reflexivity, on the other hand, seems to be an undesired property, as it produces the rankings $1 < 1$, $2 < 2$, and $3 < 3$, which in combination with the G-PCC IO $\not<$ DO would always block combinations of clitics with the same person specification. However, recall that we do deliberately excluded these cases earlier on, so reflexivity is actually unproblematic in the context of this paper.

A binary relation that is transitive and reflexive but not necessarily antisymmetric, as is the case for graph reachability, is called a *preorder*. So the space of 64 possible combinations can be narrowed down to at least those that are the result of interpreting $<$ in the G-PCC as a preorder over the set $\{1, 2, 3\}$. This reduces the number of possible PCCs to 29, which is a lot less but still too much. We need to put further restrictions on the class of possible locality graphs in order to get down to the four attested ones.

2.3 A Stronger Restriction on Rankings: Presemilattices

One way in which these four graphs differ from arbitrary preorders is that every node is related to at least one other node. A relation that does not order 2 with respect to 1 and 3 could still be a preorder, but it would not be a suitable locality metric for our purposes. This shows that we must enforce a limited amount of connectedness. Totality — the property that every node is related to every node — would be too strong a requirement, since the W-PCC and the M-PCC rely on the fact that two specific nodes are not ordered with respect to each other.

The U-PCC and W-PCC graphs share the property that they are *meet semilattices*: for any two x and y in the graph, there is a unique z such that z is reachable from both x and y and there is no z' that is distinct from z and closer to x and y (x, y, and z need not be distinct). Similarly, the U-PCC and M-PCC graphs are *join semilattices*: for any two x and y, there is a unique z such that x and y is reachable from z and there is no z' that is distinct from z and closer to x and y (once again x, y, and z need not be distinct). The S-PCC is not a semilattice for if we pick 1 for x and 2 for y, then z can be both 1 or 2, so it is not unique. However, the S-PCC is still a *presemilattice* (Plummer & Pollard 2012), and since every semilattice is a presemilattice, all four graphs form semilattices.

Definition 1. *Let $\sqsubseteq$ be a preorder on some set A. Given some subset B of A, $a \in A$ is a lower bound of B iff $a \sqsubseteq b$ for every $b \in B$. Similarly, $a \in A$ is an upper bound of B iff $b \sqsubseteq a$ for every $b \in B$. A lower bound a of B is a greatest lower bound iff B has no lower bound c such that $a \sqsubseteq c$ but not $c \sqsubseteq a$. And an upper bound a of B is a least upper bound iff B has no upper bound c such that $c \sqsubseteq a$ but not $a \sqsubseteq c$.*

Definition 2. *Let $\sqsubseteq$ be a preorder on some set A. A binary operation $\sqcap$ ($\sqcup$) is called a* meet *(*join*) operation if for all $a, b \in A$, $a \sqcap b$ ($a \sqcup b$) is a greater lower bound (least upper bound) of $\{a, b\}$ with respect to $\sqsubseteq$; note that $a \sqcap b$ ($a \sqcup b$) need not be unique. We call $\langle A, \sqsubseteq, \sqcap \rangle$ a* meet presemilattice *(a* join presemilattice *for $\sqcup$).*

These definitions are fairly technical, but they capture exactly the limited kind of connectedness that is needed here: for any two nodes, we can find a closest point where the arrows leading from them (or to them) meet, and if there are multiple such intersection points, they must be related to each other. In sum, we can restrict the possible class of relations even further to those that define presemilattices over $\{1, 2, 3\}$.

2.4 The Last Restriction: 1 is Maximal, 3 is Minimal

Even with the restriction to presemilattices our class of graphs is still too big. This time the overgeneration is due to the fact that all three person features are treated as equals, which fails to exclude orders that are the image of one of the four intended presemilattices under some non-trivial permutation of $\{1, 2, 3\}$ (i.e. any permutation that's not the identity map). The S-PCC, W-PCC, and M-PCC are each isomorphic to 2 unattested patterns, and the U-PCC to 5. In addition we also fail to rule out the structures in Fig. 2, and variations thereof. Adding up all the variants, the restriction to presemilattices still leaves us with 19 possible graphs.

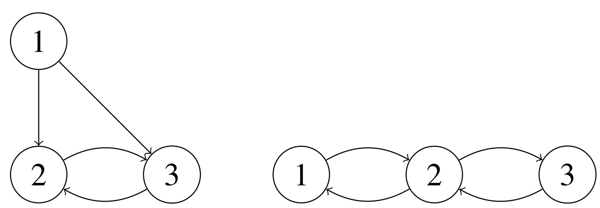

Figure 2: Two more presemilattices over the set of person features

Two natural conditions on the distribution of 1 and 3 suffice to patch this loop-hole and finally give us a full characterization of the relevant locality metrics.

Top : For all x, $1 < x$ implies $x < 1$.

Bottom : There is no x such that $x < 3$.

In other words, 3 is always the unique minimal element, and 1 always a maximal element. While these conditions are very specific, there is some independent evidence for them. In binding theory, first and second person reflexives are less restricted in their distribution than third person reflexives, which shows that third person enjoys special status. The special status of first person can be seen in cases of resolved agreement between finite verbs and coordinated subjects, the person inflection on the verb must be first person if one of the conjuncts is first person.

2.5 Interim Summary

Starting with the observation that the definition of the U-PCC can be generalized to encompass all versions of the PCC by varying the locality relations between person features, we have narrowed down the class of locality hierarchies to the attested four by imposing algebraic restrictions on their graph-theoretic representation. The class of attested PCCs is thus characterized by two components:

1. The G-PCC IO $\nless$ DO,

2. the definition of $<$ in terms of graphs over the set $\{1, 2, 3\}$ of person features that

   - are presemilattices, and
   - satisfy both *Top* and *Bottom*.

3 Deriving the Presemilattices

3.1 Rediscovering Zwicky's Person Hierarchy

While presemilattices are a fairly natural algebraic structure, they are still a highly specialized type of object, which makes one wonder why they should play a major role in limiting the PCC. A strong trend in the syntactic literature is the reduction of the PCC's pecularities to more basic assumptions about the relative complexity of person features. Nevins (2007), for example, decomposes person features into bundles of more primitive features, and then derives various facts from the differences in complexity between these bundles. On an algebraic level we may interpret this as defining a complexity metric over feature bundles that results in a basic ranking of these tree features. A promising next step, then, would be to identify a more basic ranking of person features from which the four graphs can be derived. Once this ranking is in place, we may try to reduce it to feature-geometric properties in the spirit of Nevins (2007).

In technically more precise terms, what we need is some graph over the person features that can only be mapped to the four PCC locality hierarchies given a few structural restrictions on what constitutes a licit mapping. Let's call this graph the *universal base*. Our goal is to find a universal base such that the four PCC graphs can be derived by it by only marginally changing the structure while preserving all its essential properties.

Whatever the universal base may look like, we can tell that it must establish an asymmetry between first and second person, for two reasons. First, the U-PCC is an attested PCC, but the variant where the position of 1 and 2 is switched is not. Similarly, the M-PCC is attested, but if we exchange the positions of 1 and 2 we obtain an unknown PCC. If 1 and 2 had equal standing in the universal base, then a mapping that turns the universal base into the U-PCC or the M-PCC could also turn it into the unattested variants (unless we explicitly block this in the mapping, which is unappealing). Using almost the same argument, we can also infer that there must be an asymmetry between second and third person because the variant of the U-PCC where 2 and 3 are switched does not exist, nor does the analogues version of the W-PCC.

These two asymmetries could be summarized as $2 < 1$ and $3 < 2$. So the universal base must be some structure over the set of person features that enforces these asymmetries. But that is exactly the structure of the U-PCC (modulo the transitivity arrow representing $3 < 1$). In other words, if there is a universal base for the four PCC graphs, then it must be the U-PCC graph (or a trivial alternative, e.g. one with the direction of the arrows swapped). The true appeal of this result stems from the close connection between the U-PCC graph and the well-known person hierarchy of Zwicky (1977). Since it is legitimate to equate the two, we may conclude that Zwicky's person hierarchy forms the universal base for the PCC graphs.

3.2 PCC-Graphs via Structure-Preserving Maps

We could safely identify what the universal base must look like without making any assumptions about what the licit mappings look like, except that they are structural in nature and thus do not pay attention to whether a node is labeled 1, 2, or 3. But if we want to characterize the class of PCC-graphs as the image of the universal base under a class of mappings, we have to precisely define this class of mappings. We can approach this issue by asking which properties of the U-PCC graph hold in all graphs, and which ones do not.

Suppose that we add the implicit arrow from the U-PCC graph such that there is also an arrow from 1 to 3. Then in order to get the S-PCC, we have to add an arrow between 1 and 2. The W-PCC and the M-PCC, on the other hand, are obtained by removing an arrow, either between 1 and 2, or between 2 and 3. If we do not change anything, we get the U-PCC, of course. At the same time, we cannot add an arrow between 2 and 3, and removal of arrows is limited to at most one of the two options listed before. While all of this seems pretty *ad hoc*, it actually amounts to the requirement to preserve three properties of the universal base: *Top*, *Bottom*, and *Non-Insularity*.

Non-Insularity For every x, there is y distinct from x such that $x < y$ or $y < x$.

Top and *Bottom* can also be paraphrase as conditions on the mapping.

Maximality For every mapping h and node x, if there is no y such that $x < y$, then for every y it holds that $h(x) < h(y)$ only if $h(y) < h(x)$.

Minimality For every mapping h and node x, if there is no y such that $y < x$, then there is no y such that $h(y) < h(x)$.

This has the advantage that *Top* and *Bottom* become epiphenomenal: the mappings are required to preserve maximality and minimality in a specific fashion, and 1 and 3 just happen to be maximal and minimal elements in the universal base. A different base, say, for gender features, could thus produce constraints that differ significantly from *Top* and *Bottom*.

In sum, our original characterization of the class of PCCs now involves three components:

1. The G-PCC IO $\not<$ DO,
2. the universal base $3 < 2 < 1$,
3. the class of mappings that preserve
   - non-insularity, and
   - maximality, and
   - minimality.

Although the number of components has increased, this is still preferable since the individual components are less stipulative. The universal base is very natural from a linguistic perspective, and the preservation requirements are also simple and intuitive.

3.3 Feature Geometries and the Universal Base

Zwicky's person hierarchy makes for a very plausible universal base, but it is specific to person features. If we want to push the algebraic approach to morphosyntactic dependencies beyond the realm of the PCC, it would be preferable to have a general procedure for computing universal bases for each area. Just like the PCC graphs can be reduced to a single universal base plus some preservation requirements on structural maps, the universal bases for person, number, gender and so on should be reducible to a single underlying structure with certain properties that must be maintained. In the syntactic literature, this role is filled by feature decomposition and feature geometries.

Following Nevins (2007), for example, we may posit that person features are bundles of binary-valued features.

Person	**Specification**
1	[+author, +participant]
2	[−author, +participant]
3	[−author, −participant]

The person hierarchy then can be derived by a complexity metric that posits $x < y$ iff x has fewer positive features than y. Unfortunately, though, this is just one of many different ways of deriving the person hierarchy. For one thing, any feature system will derive the hierarchy where 1 has more positive features than 2, which in turn has more than 3. It does not matter whether person features are bundles of two features or fifty. In fact, they could even have infinitely many features as long as only finitely many can have a positive value. Similarly, we could use a privative feature system and only count the number of features per bundle. Without restrictions on what constitutes a licit complexity metric, there is no way to posit a unique underlying structure like we did for the universal base.

Things are not completely hopeless, however. For one thing, evidence from other areas of morphosyntax can favor certain feature systems and complexity metrics over others. Second, some feature geometries have the advantage of needing less elaborate complexity metrics. In the feature system of Harley & Ritter (2002a,b) for instance, first and second person are on equal footing. They have the same number of features — first person is the bundle {referent, participant, author}, whereas second person is {referent, participant, addressee}. Third person is simply {referent}. In addition, first and second person have identical tree-geometries:

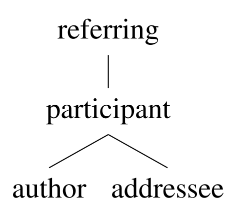

The only way, then, to obtain a universal base with $2 < 1$ is to use a complexity metric that assigns greater importance to *author* than to *addressee*. This is unnecessary for Nevins' feature system.

4 Generalizing the Generalized PCC

So far the status of the G-PCC has gone entirely unquestioned, in particular why it should take the form IO $\not<$ DO, rather than IO $<$ DO, DO $\not<$ IO, or

DO < IO. Surprisingly, the specific choice of constraint is not as important as one would expect.

Table 2 shows how the PCC data patterns in Tab. 1 can be reinterpreted as specific orderings of the person features depending on what kind of locality relation the G-PCC enforces between the direct and the indirect object. It is readily apparent that the patterns for IO $<$ DO and DO $\nless$ IO, which are depicted in (2(a)) and (2(b)), are just the respective complements of IO $\nless$ DO and DO $<$ IO, shown in (2(d)) and (2(c)). Moreover, (2(a)) and (2(b)) differ in only two respects. The first and the third column have been switched, revealing a close algebraic relation between the S-PCC and the W-PCC. Second, the values for the M-PCC differ for $2 < 3$ and $3 < 2$. For one version of the G-PCC the two person feature are unordered for the M-PCC, for the other one they form a loop in the algebra.

We already encountered this structure in Fig. 2 as an example of a presemilattice that does not satisfy *Bottom*. If we want to allow this structure, we have to weaken *Bottom* to *Bottom*′ so that it more closely resembles *Top*. Equivalently, we could relax *Minimality*.

Bottom′ For all x, $x < 3$ implies $3 < x$.

Minimality′ For every mapping h and node x, if there is no y such that $y < x$, then for every y it holds that $h(y) < h(x)$ only if $h(x) < h(y)$.

Notice that this also brings in the other structure in Fig. 2. So the class of licit graphs now has increased from four to six.

Interestingly, this has little effect on the predicted typology. Let us refer to the presemilattices as G1, G2, …, G6 according to the order in which they occurred in this paper. So G1 is the S-PCC graph, and G5 is the left graph in Fig. 2. As shown in Tab. 3, both IO $<$ DO and DO $\nless$ IO induce all attested PCCs over these six graphs.

Both constraints also produce a new type of PCC, which I call the *Me second PCC* (M2-PCC). The M2-PCC prohibits DO clitics from combining with anything but 1IO clitics. Rather than the M-PCC's ban against 1DO with 2/3IO, then, we obtain a requirement for 1IO. To the best of my knowledge, this type of PCC is unattested.

However, the table contains two more variants of the PCC that haven't been discussed so far and that do exist, although they are not necessarily classified as PCCs. Both types are induced by the sixth graph, in which every person feature is less local than all other features. With DO < IO, this yields the *Null PCC* (N-PCC), which deems all logically possible clitic combinations grammatical. The shift to IO $\nless$ DO, on the other hand, blocks all clitic combinations indiscriminately, which is why we may refer to this as the *Indiscriminate PCC* (I-PCC). Martin Walkow (p.c.) points out that this constraint is active in languages that disallow cliticization of more than one object, such as Cairene Arabic (cf. Shlonsky 1997:207).

The other two G-PCCs, IO $<$ DO and DO $\nless$ IO produce very different typologies unless they are computed over the complements of the six graphs — i.e. the graphs where x and y are connected by an arrow iff they are not connected by an arrow in the original graph. This suggests that the definition of the G-PCC should itself be tied more closely to the graphs that encode the locality ranking. If the shape of the graphs informs which two out of the four G-PCCs may be chosen, then the G-PCC has only a marginal effect on

Order	S-PCC	U-PCC	W-PCC	M-PCC
1 < 2	$*$	$\checkmark$	$\checkmark$	$\checkmark$
1 < 3	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$
2 < 1	$*$	$*$	$\checkmark$	$*$
2 < 3	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$
3 < 1	$*$	$*$	$*$	$*$
3 < 2	$*$	$*$	$*$	$\checkmark$

(a) IO $<$ DO

Order	S-PCC	U-PCC	W-PCC	M-PCC
1 < 2	$\checkmark$	$\checkmark$	$*$	$\checkmark$
1 < 3	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$
2 < 1	$\checkmark$	$*$	$*$	$*$
2 < 3	$\checkmark$	$\checkmark$	$\checkmark$	$*$
3 < 1	$*$	$*$	$*$	$*$
3 < 2	$*$	$*$	$*$	$*$

(b) DO $\not<$ IO

Order	S-PCC	U-PCC	W-PCC	M-PCC
1 < 2	$\checkmark$	$*$	$*$	$*$
1 < 3	$*$	$*$	$*$	$*$
2 < 1	$\checkmark$	$\checkmark$	$*$	$\checkmark$
2 < 3	$*$	$*$	$*$	$*$
3 < 1	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$
3 < 2	$\checkmark$	$\checkmark$	$\checkmark$	$*$

(c) IO $\not<$ DO

Order	S-PCC	U-PCC	W-PCC	M-PCC
1 < 2	$*$	$*$	$\checkmark$	$*$
1 < 3	$*$	$*$	$*$	$*$
2 < 1	$*$	$\checkmark$	$\checkmark$	$\checkmark$
2 < 3	$*$	$*$	$*$	$\checkmark$
3 < 1	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$
3 < 2	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$

(d) DO $<$ IO

Table 2: Overview of order relations by G-PCC variants

Structure	IO $\not<$ DO	DO < IO
G1	S	W
G2	U	U
G3	W	S
G4	M	M2
G5	M2	M
G6	I	N

Table 3: Extended class of presemilattices and the PCCs they encode

the predicted typology. In fact, if the M2-PCC should turn out to exist after all, the choice of G-PCC would only have an effect on the availability of the I-PCC and the N-PCC, both of which are somewhat pathological.

Conclusion

An algebraic perspective of morphosyntax that completely abstracts away from matters of syntactic implementation offers a succinct characterization of the attested PCC typology in terms of two components: a generalized PCC that references a person locality hierarchy, and a class of graphs that represents the possible hierarchies. These graphs, in turn, are exactly those that preserve important structural properties of the universally fixed person hierarchy $3 < 2 < 1$. While there are still some open questions, in particular regarding the derivability of Zwicky's person hierarchy and the definition of the generalized PCC, the algebraic approach shows great potential. It will be interesting to see whether other aspects of morphosyntax can be analyzed along similar lines, and if so, what the algebraic commonalities will turn out to be.

References

Adger, David, & Daniel Harbour. 2007. The syntax and syncretisms of the person case constraint. *Syntax* 10.2–37.

Anagnostopoulou, Elena. 2005. Strong and weak person restrictions: A feature checking analysis. In *Clitics and Affix Combinations: Theoretical Perspectives*, ed. by Lorie Heggie & Francisco Ordoñez, 199–235. Amsterdam: John Benjamins.

Béjar, Susana, & Milan Rezac. 2009. Cyclic agree. *Linguistic Inquiry* 40.35–73.

Bobaljik, Jonathan D. 2012. *Universals in Comparative Morphology: Suppletion, Superlatives, and the Structure of Words*. Cambridge, Mass.: MIT Press.

Bonet, Eulàlia. *Morphology After Syntax: Pronominal Clitics in Romance*. MIT, Boston, MA dissertation.

——. 1994. The Person-Case Constraint: A morphological approach. In *The Morphology-Syntax Connection*, number 22 in MIT Working Papers in Linguistics, 33–52.

Corbett, Greville G. 1991. *Gender*. Cambridge, Mass.: Cambridge University Press.

Harley, Heidi, & Elizabeth Ritter. 2002a. Person and number in pronouns: A feature-geometric analysis. *Language* 78.482–526.

——, & ——. 2002b. Structuring the bundle: A universal morphosyntactic feature geometry. In *Pronouns — Grammar and Representation*, ed. by Horst J. Simon & Heike Wiese, 23–39. Amsterdam: John Benjamins.

Kayne, Richard S. 1975. *French Syntax: The Transformational Cycle*. Cambridge, Mass.: MIT Press.

Nevins, Andrew. 2007. The representation of third person and its consequences for person-case effects. *Natural Language and Linguistic Theory* 25.273–313.

——. 2011. Multiple agree with clitics: Person complementarity vs. omnivorous number. *Natural Language and Linguistic Theory* 28.939–971.

Plummer, Andrew, & Carl Pollard. 2012. Agnostic possible world semantics. In *LACL 2012*, ed. by Denis Béchet & Alexander Dikovsky, number 7351 in Lecture Notes in Computer Science, 201–212, Heidelberg. Springer.

Shlonsky, Ur. 1997. *Clause Structure and Word Order in Hebrew and Arabic*. Oxford: Oxford University Press.

Walkow, Martin. *Goals, Big and Small*. University of Massachusetts Amherst dissertation.

Zwicky, Arnold. 1977. Hierarchies of person. In *Chicago Linguistic Society*, volume 13, 714–733.