



The evolution of understanding of syntactic structure and its mathematical representation in the structuralist school: From Saussure to set theory, graphs, and algebras

Evans Gesura Mecha

evans-mecha@kisiiuniversity.ac.ke

Kisii University, Kenya

Recommended Reference: Mecha, E. G. (2026). The evolution of understanding of syntactic structure and its mathematical representation in the structuralist school: From Saussure to set theory, graphs, and algebras. *African Quarterly Social Science Review*, 3(2), 591–601. <https://doi.org/10.51867/AQSSR.3.2.50>

ABSTRACT

What kind of mathematical object is a sentence? This article traces how structuralist linguistics has answered that question, from Saussure's sign-system through Bloomfieldian and Harrisian distributionalism to the labelled tree of early generative grammar, the unordered set of Minimalist Merge, and today's graph-theoretic and algebraic reformulations. The paper adopts a conceptual-historical and critical analytical approach, examining foundational and contemporary theoretical literature on linguistic structuralism, mathematical representation, and the philosophy of scientific modelling. Its research design involves a comparative analysis of successive models of syntactic structure, focusing on their underlying mathematical assumptions, explanatory capacities, and limitations. We disentangle theses often conflated under “structuralism”: the linguistic school, mathematical structuralism, structural realism, and the structuralist theory of science, drawing on Rickart, Hellman and Shapiro, Votsis, Balzer, Sneed and Moulines, and Nefdt's model-based reading of grammar. We then re-examine linearization, the deep/surface distinction, binarity, endocentricity, Merge, and Hopf-algebraic and graph-theoretic reformulations, arguing that this history isolates which relational properties of language are genuinely forced and which are formal artefacts of a chosen idiom. Findings reveal that c-command remains a relatively stable relational property across changes in syntactic representation, whereas the deep-surface structure distinction illustrates how theoretical commitments may arise from the formalism adopted. The paper contributes a critical framework for distinguishing descriptive adequacy, derivational economy, and formal overhead in evaluating competing mathematical representations of syntactic structure.

Keywords: C-command, Linearization, Mathematical Structuralism, Merge, Model-Based Linguistics, Multidominance, Structural Realism, Structuralism, Syntactic Structure, X-Bar Theory

I. INTRODUCTION

Every syntactic theory commits itself, usually without saying so, to a mathematical ontology. To say that a sentence “has structure” already presupposes a formal object capable of bearing that structure: a system of oppositions, a string, a tree, a set, a graph, an algebra. The choice among these is rarely presented as a choice at all. It is inherited, together with a battery of empirical successes and failures, from whichever theory made it first. The history of syntax across the structuralist lineage, from Saussure onward, is in large part a history of this choice being revisited under pressure from specific anomalies that the reigning mathematical idiom could not comfortably absorb.

This article reconstructs that history, but only after pausing on a prior question the syntax literature tends to skip: what, exactly, is meant by “structure”, and by “structuralism”? These terms cover at least four distinct theses that are historically entangled but logically independent. There is structuralism as a linguistic school, the Saussurean and post-Saussurean view that a language is a system of interdependent values. There is mathematical structuralism, a position in the philosophy of mathematics associated with Benacerraf (1965), Resnik (1997), Shapiro (1997), and, in systematic overview, Hellman and Shapiro (2019), holding that mathematical objects have no nature beyond their place in a relational pattern. There is structural realism, a position in the philosophy of science associated with Worrall (1989) and developed by Ladyman and Ross (2007) and Votsis (2011), holding that what survives theory change in a mature science is relational structure rather than the intrinsic nature of the objects instantiating it. And there is the structuralist theory of science of Balzer, Sneed, and Moulines (Balzer & Moulines 1996; Balzer, Sneed, & Moulines 2000), a formal metatheoretic apparatus for reconstructing empirical theories as classes of models. Nefdt (2016, 2019, 2021) has argued that generative grammar is best interpreted through structural realism and, more broadly, through the philosophy of scientific modelling, rather than through either a naive Platonism about grammars or a purely instrumentalist reading of them; and Meier (2012) has shown, through a direct formal reconstruction of Bloomfield's

own theory, that even the informal, taxonomic wing of American structuralism already had an implicit model-theoretic structure waiting to be made explicit.

Section 2 develops this conceptual groundwork, including the sense in which a grammar is a scientific model rather than either a bare formalism or a direct transcript of psychological reality. Section 3 traces the genealogy of syntactic representation from Saussure through Bloomfield, Harris, and Wells to Chomsky's labelled tree and onward to Minimalist Merge, noting an early and largely forgotten algebraic alternative, Lambek (1958) categorial grammar. Section 4 examines the flat-versus-hierarchical controversy and Kayne (1994) Linear Correspondence Axiom. Section 5 argues that the deep/surface structure distinction was an artefact of a particular rewriting formalism rather than an empirical discovery. Sections 6 through 8 examine the graph-theoretic, set-theoretic, and algebraic idioms in which Merge and its predecessors have been formalized, and Section 9 gathers the graph-theoretic critique of set-theoretic Merge advanced by Krivochen (2023) and Zyman (2024). Section 10 returns to structural realism to ask what has survived across these reformulations, and Section 11 offers a critical synthesis before the conclusion.

1.1 Research Objectives

To critically examine the evolution of mathematical representations of syntactic structure, identifying the relational properties that remain stable across theoretical models and those that arise from the formalism adopted.

II. LITERATURE REVIEW

2.1 What is a structure? Grammars as scientific models

Rickart (1995), a mathematician rather than a linguist or a philosopher, offers a usefully deflationary starting point: a structure, in the most general mathematical sense, is a set of objects together with a set of relations and operations defined on them, and “structuralism” in any discipline names the strategy of analysing a subject matter by specifying such a structure rather than by cataloguing the intrinsic properties of its elements one by one. On this broad reading, Saussurean linguistics, Bloomfieldian distributionalism, X-bar theory, and Merge-based Minimalism are all structuralist in exactly the same sense: each specifies a domain of objects (signs, forms, phrase markers, syntactic objects) and a set of relations or operations over that domain (value, distribution, dominance, Merge), and takes the relations, not the objects' intrinsic nature, to be doing the explanatory work. Rickart's own discussion extends this treatment to linguistics as one of several fields, alongside mathematics and biology, in which the structuralist method proves indispensable.

Mathematical structuralism proper, as surveyed by Hellman and Shapiro (2019), is narrower and more specifically ontological: it holds that mathematical objects such as numbers have no nature beyond the relations they bear to other objects within a pattern, so that asking whether “the number 2” is really a set, an ordinal, or a category-theoretic object is a confused question. When a syntactician says that a syntactic object “is” a set, a graph, or an algebra, they are implicitly taking a position in this same debate, whether or not they intend to: on a structuralist reading of mathematics, the fact that Merge can be equally well described as forming an unordered set, a labelled node in a rooted tree, or a generator in a free magma is not itself troubling, because none of these descriptions is more truly what a syntactic object “is” than any other, provided each preserves the relevant relational pattern.

Structural realism, by contrast, is a thesis in the philosophy of science: Worrall (1989) proposed it to explain how successive, mutually incompatible physical theories can each be approximately true, by holding that what is preserved across theory change is the mathematical structure of the equations rather than the ontology of the entities the equations were originally taken to describe. Ladyman and Ross (2007) radicalized this into ontic structural realism, on which relational structure is not merely what we can know about the world but what the world fundamentally is. Votsis (2011) argues for a more measured, epistemically framed version of the same continuity claim, cataloguing the historical cases in which structural continuity does and does not hold and cautioning against treating every case of theory change as automatically vindicating the structuralist diagnosis. Nefdt (2021) imports this apparatus directly into generative linguistics, arguing that a structural realist reading of grammars, on which what is real is the pattern of syntactic relations a grammar describes, not the reified phrase markers or feature bundles any particular formalism uses to describe it, makes better sense of theory change in syntax than either an uncompromising Platonism about grammatical objects or an eliminativist instrumentalism.

Nefdt's (2016, 2019) companion methodological argument brings us to the question of mathematical representation and modelling proper. Grammars, on Nefdt's account, are best understood neither as directly represented psychological states nor as purely formal theories on the model of arithmetic, but as scientific models occupying an intermediate “grade” of mathematical involvement, comparable to the models used elsewhere in the natural sciences. This reading connects directly with the general philosophy of scientific modelling: Godfrey-Smith (2006) characterizes model-based science as a strategy in which the theorist gains understanding of a complex real-world system via understanding of a simpler, hypothetical system that resembles it in the relevant respects, rather than

via a single, exhaustive theory stated in the object language of the domain itself. A labelled tree, an unordered set, or a Hopf algebra is, on this reading, exactly such a hypothetical system: none of them is offered as a literal transcript of what happens in a speaker's head, and each is judged by how well its structure resembles the structure of the syntactic dependencies it is built to capture, not by whether its constituent formal objects (sets, nodes, algebra elements) literally exist inside the language faculty. This is, in fact, a much older idea in linguistics than the philosophy-of-science literature that later made it explicit. Harris's (1951) discovery procedures for distributional analysis already functioned as a model-building methodology in Godfrey-Smith's sense: the analyst constructs a simplified formal object (a set of substitution classes, later a set of transformations relating sentence types) that resembles the corpus of attested utterances closely enough, in the relevant distributional respects, to support generalization beyond the corpus actually observed, without claiming that the formal apparatus itself is psychologically real for the speaker. Read this way, the shift from Harris's substitution classes to Chomsky's phrase markers to Minimalist Merge is not simply a shift in mathematical idiom for its own sake, but a sequence of competing models judged, as Godfrey-Smith's framework would predict, by their resemblance to an independently specified target: the native speaker's judgements of constituency, dependency, and grammaticality, not by the intrinsic elegance of set theory, graph theory, or algebra as such.

With this apparatus in hand, we can state precisely what is at stake in the rest of this article. Whether syntactic structure is “really” a tree, a set, a graph, or an algebra is not, on the mathematical-structuralist reading of the second paragraph above, a fully well-posed question unless further constrained: it must be read either as a question about which formal idiom most economically models the relational pattern linguists have independently discovered (a question of descriptive and derivational adequacy, addressed in Sections 3, 4, and 6–9), or as a question about which relational pattern, among several formally distinct candidates, actually survives across successive syntactic theories in the sense structural realism intends (a question we return to in Section 10).

III. METHODOLOGY

The study adopts a conceptual-historical methodology to trace the evolution of mathematical representations of syntactic structure, from Saussurean structuralism and distributional approaches to generative grammar, Minimalist Merge, and contemporary graph-theoretic and algebraic models. It draws on foundational and contemporary theoretical literature in linguistics, mathematics, and the philosophy of science to examine how different approaches conceptualise and formalise syntactic relations.

The study employs a comparative and critical analytical design, evaluating the assumptions, explanatory capacities, and limitations of the selected models. The analysis focuses on key issues, including c-command, linearization, deep and surface structure, binarity, endocentricity, and Merge. These are examined to distinguish relational properties of syntax that persist across models from those that depend on particular mathematical formalisms, while assessing the descriptive adequacy, derivational economy, and formal overhead of competing representation

IV. FINDINGS & DISCUSSIONS

4.1 From sign to tree: A genealogy of syntactic structure

Saussure's posthumously assembled *Cours de linguistique générale* (Saussure 1959 [1916]) established the founding move of linguistic structuralism: a sign is not a label attached to a pre-existing thing but an associative bond between a signifier and a signified, and the value of any sign is fixed not by reference to the world but by its place within the whole system of oppositions constituting a language at a given moment (*langue*, as opposed to the individual acts of *parole*). Syntax, in the Saussurean system, was addressed only obliquely, through the notion of syntagmatic relations, the linear, in praesentia combination of signs in an utterance, contrasted with the associative, in absentia relations linking a sign to the paradigm of forms it could have replaced. Saussure did not offer, and did not attempt, a generative or recursive account of how syntagms are built; the structuralist method, at this stage, was a method for describing a system of values, not a method for generating an unbounded set of well-formed strings.

American structural linguistics, running from Bloomfield (1926, 1933) through Harris (1951), inherited the emphasis on system-internal relations but recast it in a more explicitly empiricist idiom. Bloomfield's (1926) postulates defined linguistic forms recursively in terms of an underlying set of utterances, a partial ordering of “free” and “bound” forms, and a notion of formal and semantic identity across occurrences. This apparatus, as Meier's (2012) reconstruction shows, is a genuine, if informally stated, model-theoretic structure in the sense of Section 2. Constituency itself was given its first systematic formal treatment by Wells (1947), whose “Immediate Constituents” treated a sentence's hierarchical decomposition into successively smaller constituents as a mathematical object in its own right, subject to explicit well-formedness conditions rather than to unreflective bracketing. Harris (1951)

extended distributional analysis from morphemes to whole sentences and, in later work collected in Harris and Hiz (1981), introduced the term “transformation” to describe systematic correspondences between sentence patterns, though for Harris, as noted in Section 2, transformations remained a device for normalizing data and establishing equivalence classes among attested strings: a model of the corpus, not a generative mechanism for producing new sentences. Carnie's (2008) survey of constituent structure situates these early American approaches alongside their competitors: dependency grammar, categorial grammar, relational grammar, and the various phrase-structure-based frameworks that followed. It treats constituency itself, rather than any one theory of it, as the object of comparative formal study.

A largely independent algebraic idiom for representing sentence structure was proposed by the mathematician Lambek (1958) in “The mathematics of sentence structure”. Lambek's categorial grammar assigns each lexical item a type built from a small set of basic categories using two directional slash operators, so that an intransitive verb, for instance, might be assigned the type of “an expression that combines with a noun phrase to its left to form a sentence”, and a transitive verb the type of “an expression that combines with a noun phrase to its right to form an intransitive verb”. Syntactic combination is then reduced to type cancellation, governed by algebraic laws (residuation) borrowed from the theory of residuated semigroups. Lambek's proposal did not displace phrase-structure grammar in mainstream syntactic theorizing, but it established, a full generation before Merge, that an algebraic rather than a graph-theoretic idiom could in principle bear the entire explanatory weight that trees were later asked to carry. This is a precedent worth recalling when we turn, in Section 8, to the algebraic reformulations of Merge itself.

Chomsky's (1957) Syntactic Structures reorganized the field around a different formal commitment: a phrase-structure grammar, augmented by transformations that generates rather than merely classifies the sentences of a language, with a structural description taking the form of a labelled, ordered tree. Chomsky (1965) refined this into the Aspects model's distinction between deep and surface structure, and Chomsky (1970) introduced X-bar theory, imposing a uniform endocentric template on phrasal projection across categories. The trajectory from there to Minimalism, sketched further in Sections 5 through 8, is a trajectory of successive subtraction: Government-and-Binding theory (Chomsky 1981, 1986) retained the labelled tree while deriving more of its properties from independent modules; Bare Phrase Structure (Chomsky 1994) eliminated the phrase-structure rules themselves; and the Minimalist Program (Chomsky 1995) replaced rule-governed tree construction with a single recursive set-forming operation, Merge, whose properties Chomsky and colleagues (2023) have since collected under the label of the Strong Minimalist Thesis.

4.2 Flat versus hierarchical structure: Linearization illustrated

A recurring undercurrent in this history is the suspicion that hierarchical structure is not primitive but is imposed on an underlyingly flat, linear object, the utterance as actually produced and perceived. Saussure's own syntagmatic relations, after all, were explicitly linear (in praesentia) rather than hierarchical, and Bloomfieldian immediate-constituent analysis, whatever its hierarchical trappings, was arrived at by successively partitioning a linear string rather than by generating structure top-down. The question sharpens once hierarchy is taken seriously as a graph-theoretic commitment, because a tree, unlike a string, carries no intrinsic linear order among sisters: something must explain why trees are pronounced as strings at all.

Kayne's (1994) Linear Correspondence Axiom (LCA) is the most influential answer. Kayne proposes that linear precedence at the phonological interface is derived, not stipulated, from asymmetric c-command in the hierarchical structure: if a non-terminal node A asymmetrically c-commands a non-terminal node B, every terminal dominated by A precedes every terminal dominated by B. A simple illustration makes the point concrete. Consider the tree for “John read the book”, with a VP node dominating V (read) and a DP dominating D (the) and NP (book): because V asymmetrically c-commands the DP, and D asymmetrically c-commands N within it, the LCA correctly predicts the total order read – the – book without any separate statement of linear order being written into the phrase-structure rules at all. Fox and Pesetsky (2005) later relaxed the strict correspondence between one hierarchical structure and one linearization by proposing cyclic linearization, in which order statements are computed and preserved phase by phase, weakening Kayne's claim that order is a single global readout while preserving the core idea that order is derivative of hierarchy.

The flat-structure alternative resurfaces wherever the hierarchical account requires machinery that looks ad hoc. Discontinuous constituents, right-node raising, parentheticals, certain coordination and extraposition facts, are the classic testing ground. McCawley (1982) argued that a subset of these phenomena is best modelled with node-sharing (what would later be called multidominance) rather than with additional movement or ellipsis layered onto a strictly tree-shaped representation. Blevins's (1990) dissertation develops this line systematically on the basis of Germanic verb clusters and English right-node raising, arguing that syntactic representations must in general permit a single constituent to be dominated by more than one mother, that the representation is a directed acyclic graph (DAG) that happens to be tree-shaped in the unmarked case. What this literature shows, we suggest, is that “flat versus

hierarchical” is a false dichotomy generated by an impoverished choice of mathematical object: a strictly linear string cannot state c-command-sensitive generalizations without covert appeal to hierarchy, while a strictly tree-shaped graph cannot state discontinuity facts without covert appeal to something graph-theoretically richer than a tree. The DAG is the more general object of which both the flat string (a DAG with no branching) and the classical tree (a DAG restricted to a single mother per node) are special cases.

4.3 The failure of the deep/surface structure distinction

The Aspects model (Chomsky 1965) distinguished deep structure, the level at which lexical insertion and thematic relations were represented prior to transformation, from surface structure, the level feeding phonological interpretation after transformations had applied. The distinction did real theoretical work for roughly two decades, allowing meaning-preserving transformations to be stated as structure-changing operations relating two representations of one sentence without requiring the semantic component to parse a single, maximally complex representation directly.

The distinction did not survive, we argue, primarily because the architecture that motivated it was abandoned, not because deep structure was independently falsified by a single decisive dataset. Government-and-Binding theory had already hollowed out its theoretical role by relocating scope and binding relations to a separate level, LF, that D-structure did not directly feed (Chomsky 1981, 1986). Once Merge is adopted as the sole structure-building operation, the interpretive requirement that thematic roles be assigned as soon as an argument is merged removes the need for a designated pre-transformational checkpoint, and once movement is reanalyzed as Internal Merge, a special case of the very operation that builds new structure, there is no natural point in the derivation at which to freeze a “surface” level either, since Spell-Out can apply at multiple phases rather than once (Chomsky 2000, 2013). Deep structure and surface structure, on this reading, were not natural kinds discovered by syntactic investigation but artefacts of the string-rewriting-plus-transformation formalism inherited from Chomsky (1957); when that formalism was replaced by set-theoretic Merge, the levels it had required disappeared with it, without a corresponding loss of empirical coverage. This is itself evidence that the levels had been formal overhead rather than substantive theoretical content.

4.4 Graph-theoretic representation: Binarity, X-bar, and the retreat from endocentricity

Formally, a syntactic tree is a rooted, ordered, labelled graph in which every non-terminal node dominates exactly the nodes that constitute its immediate constituents. Two further restrictions have dominated the generative tradition, both graph-theoretic in character: binarity and endocentricity. The binarity principle holds that every non-terminal node has exactly two daughters; Kayne (1994) derived this from the LCA itself, since ternary branching would leave the relative order of two sisters, neither c-commanding the other, undetermined, so that the LCA would fail to totally order the terminal string. Binarity, on this view, is a theorem of antisymmetric linearization rather than an independent stipulation.

Endocentricity, the requirement that every phrase project from a lexical head, formalized in X-bar theory as the X'/XP template (Chomsky 1970; Jackendoff 1977), fared worse. X-bar theory factored out genuine cross-categorical generalizations, but nothing in the theory explained why projection should be endocentric in the first place, and Bare Phrase Structure (Chomsky 1994) was designed to derive rather than stipulate this from the lexical content of Merge's inputs. Once Merge is taken as primitive, endocentricity becomes an empirical question about labelling rather than an architectural axiom: if Merge simply forms the set $\{\alpha, \beta\}$, nothing in the operation itself designates either α or β as the projecting head. Chomsky (2013) proposes a minimal search procedure that recovers endocentricity in the canonical head-complement case but finds no unique label when α and β are both phrases, unless independent operations resolve the symmetry. Collins (2002) eliminated labels altogether, arguing that all the empirical work labels had been assigned to do could be stated directly over sets of syntactic objects; Narita (2014) develops a projection-free architecture in the same spirit, treating endocentricity, where observed, as a derived generalization about interface interpretability rather than a primitive requirement on Merge's output. The historical shift is exactly the reverse of what X-bar theory presupposed: exocentric, unlabelled Merge is now typically taken as the null hypothesis, and endocentric projection the special case requiring explanation.

C-command, defined representationally by Reinhart (1976) over a static, already-built phrase marker as the structural condition regulating anaphoric dependencies, has proven the most durable graph-theoretic relation across these reformulations. Once Merge is adopted as the sole structure-building operation, c-command admits a derivational reformulation on which the relation is established at the moment two syntactic objects are merged, rather than read off a completed tree after the fact. This is a shift in ontological status, from a relation holding of a static structure to one constituted by a dynamic process, though its empirical predictions are left essentially intact. As we argue in Section 10, this durability of c-command across representational change is itself a useful diagnostic for distinguishing genuine relational structure from formal overhead.

4.5 Set-theoretic representation: Merge, model theory, and the class-nominalist turn

Merge is standardly defined, following Chomsky (1995, 2000), as an operation that takes two syntactic objects α and β and forms a new syntactic object $K = \{\alpha, \beta\}$, or, on labelled variants, $\{\gamma, \{\alpha, \beta\}\}$ where γ is the label. Two properties of this definition matter for what follows. First, Merge is defined over unordered sets: as a set, $\{\alpha, \beta\}$ is identical to $\{\beta, \alpha\}$, so nothing in the definition encodes which of the two precedes the other, or, on unlabelled variants, which of the two projects. Second, Merge is structure-building but interface-blind: it does not itself determine how its output is pronounced or interpreted, which is why linearization must be supplied by an independent axiom such as Kayne's (1994) LCA (Section 4). Model-theoretic treatments of grammar formalism, exemplified by Stabler's (1997) Minimalist Grammars, make this set-theoretic commitment fully explicit by formalizing Merge and Move directly as operations on labelled trees drawn from a finite lexicon, thereby giving Merge-based Minimalism the kind of formal-language-theoretic characterization phrase-structure grammars had long possessed.

A worked example makes the set-theoretic commitment concrete. Take the clause the cat chased the mouse. Merge first combines the determiner the with the noun cat, giving the set $\{\text{the}, \text{cat}\}$; call this K1. A second application of Merge combines the with mouse, giving $K2 = \{\text{the}, \text{mouse}\}$. The verb chased is then merged with K2 to form $K3 = \{\text{chased}, \{\text{the}, \text{mouse}\}\}$, and K1 is finally merged with K3 to form $K4 = \{\{\text{the}, \text{cat}\}, \{\text{chased}, \{\text{the}, \text{mouse}\}\}\}$, the syntactic object corresponding to the whole clause. Nothing in K4 as a bare set records that cat precedes chased, or that chased precedes mouse: that information is supplied only once the LCA maps asymmetric c-command onto linear order at Spell-Out. The labelling problem of the next paragraph is equally visible in this derivation: K1 and K2 are unproblematic, since the is a functional head that plausibly projects over its lexical sister, giving a determiner phrase in each case; K3 is likewise unproblematic, since chased is a lexical head projecting over its complement, giving a verb phrase. K4, however, merges two phrases, K1 and K3, neither of which is a lexical item, so minimal search finds no head to project a label directly. Standard Minimalist practice resolves this by appeal to an independent probe, typically a tense or agreement feature shared between the subject and the verbal projection, precisely the kind of externally supplied resolution the class-nominalist reading of Section 7's next paragraph treats as unproblematic in principle but as evidence that the set-theoretic object alone underdetermines category membership.

The set-theoretic reconstruction of syntactic structure raises a question the earlier graph-theoretic tradition could defer: what, exactly, is a syntactic category, if it is not a node label stipulated by rule but a derived property of an unlabelled set? Lieb's (1992) programme of "New Structuralism", developed within Integrational Linguistics, offers a class-nominalist answer: a grammatical category such as 'noun' or 'verb phrase' is not a further entity possessed by the items that instantiate it, over and above the class of those items itself, but simply is the class, with category membership exhausted by class membership and no additional universal doing further explanatory work. Applied to Merge-based syntax, this yields a deflationary reading of the labelling problem: to ask what $\{\alpha, \beta\}$'s label is, on a class-nominalist reading, is to ask which independently specifiable class of syntactic objects $\{\alpha, \beta\}$ belongs to, not which further, ontologically distinct labelling-entity has been attached to the set. Minimal search (Chomsky 2013) is then not discovering a hidden label that Merge's set-theoretic definition had merely failed to display, but constructing, post hoc, a class-membership verdict from the featural properties of α and β . Lieb's New Structuralism was developed as a general programme for the foundations of linguistic theory rather than as a commentary on Minimalist Merge specifically, which postdates much of that programme, so the application sketched here is a reconstruction rather than an exegesis.

It is worth being explicit about why Merge produces unordered syntactic objects at all, since this is often treated as a brute stipulation rather than a theoretically motivated choice. The Strong Minimalist Thesis, that language is an optimal solution to interface conditions imposed by the sensorimotor and conceptual-intentional systems (Chomsky 2004; Chomsky et al. 2023), places the burden of proof on any property of Merge not independently forced by third-factor considerations. Linear order is plausibly not required by the conceptual-intentional interface, which appears largely insensitive to the order in which arguments are pronounced; it is required only by the sensorimotor interface, which must externalize a hierarchical object as a temporally ordered signal. If order is imposed by externalization alone, the null hypothesis is that Merge should not encode order at all, hence the choice of an unordered set rather than an ordered pair as Merge's output.

4.6 Algebraic representation: Categorical grammar, Hopf algebras

A third mathematical idiom treats the recursive structure-building operation as an operation within an explicitly defined algebraic structure, a residuated semigroup, a Hopf algebra rather than as an operation whose only formal property is that it forms unordered pairs. Lambek's (1958) categorial grammar, introduced in Section 3, is the earliest instance: syntactic combination there just is algebraic cancellation of directional types, with no separate tree-building operation required at all. A transitive sentence such as "John reads books" is licensed because the type of "reads" cancels against a noun phrase to its right and then against a noun phrase to its left, yielding a sentence directly from the algebra of the type system.



More recent algebraic reformulations of Merge itself have been motivated by a dissatisfaction with the bare set-theoretic picture on at least two grounds. First, an algebraic reformulation can make explicit formal properties of Merge, associativity, the existence of an identity, the behaviour of Merge under composition with Internal Merge, that are left implicit when Merge is characterized only as “form the set $\{\alpha, \beta\}$ ”. Second, Hopf algebras come equipped with a coproduct, an operation dual to multiplication that decomposes a composite object into an ordered sum of ways of splitting it into sub-objects, previously exploited in mathematical physics (renormalization in quantum field theory) to model exactly the kind of nested, recursively decomposable combinatorial objects that Merge-generated syntactic objects also are. Marcolli, Berwick, and Chomsky (2023a) compare Stabler's (1997) computational Minimalism with Chomsky's newer formulation of Merge from this point of view, showing that External Merge in Stabler's system is describable as a partially defined operated algebra while Internal Merge determines a system of coideals of the Loday–Ronco Hopf algebra, whereas the newer formulation of Minimalism carries a single Hopf-algebra structure in which Internal and External Merge arise from the same underlying operation. This is, on their argument, a mathematical advantage for the newer formulation. Marcolli, Chomsky, and Berwick (2025) develop the full formalization in book-length form, proving that Merge can be described as a particular kind of highly structured algebra, a companion line of work also pursued in Marcolli, Berwick, and Chomsky's (2023b) algebraic treatment of the syntax-semantics interface itself.

4.7 Multidominance and the graph-theoretic critique of set-theoretic Merge

A distinct and more directly oppositional line of argument holds that the right response to the limitations of unordered set-formation is not to add algebraic structure on top of it but to abandon sets in favour of graphs from the outset. De Vries (2009) develops multidominance theory within an otherwise standard Minimalist framework, arguing that certain constructions (parasitic gaps, right-node raising, some cases of coordination) are best represented as directed acyclic graphs in which a single constituent has more than one mother, rather than as trees related by covert movement or ellipsis.

Krivochen (2023) presses a sharper, more general version of this critique, focused not on any single construction but on the operation of Minimal Search itself, which Minimalist theory takes to be responsible for locating the closest goal for Agree, for determining labels, and for a range of other core computations. Krivochen argues that, given the Minimalist assumption that structure building consists of unordered set-formation, Minimal Search cannot in general be defined as a well-behaved search algorithm: certain configurations, notably $\{XP, YP\}$ and $\{X, Y\}$, are set-theoretically ambiguous between structural descriptions that ought to be distinguished, because sets encode no information about linear order, directionality of selection, or which element is closer to a given search target. If, instead, structure building is reconceived in terms of directed graphs, with argument-taking relations represented as arcs rather than as set membership, Minimal Search can be defined unambiguously as a graph-search algorithm operating over an ordered set of arcs. Zyman (2024) reaches a strikingly similar conclusion from a more conservative, definitionally focused starting point: evaluating two recent set-theoretic definitions of Merge on their own terms, he argues that both face conceptual problems inherent to any set-theoretic definition, and proposes instead that Merge is graph-theoretic rather than set-theoretic in nature, operating directly on (bare-phrase-structure-compliant) phrase-structure trees.

Returning to the cat chased the mouse from Section 7 makes the contrast tangible. A graph-theoretic rendering of the same clause does not build an unordered set at each step but adds directed arcs to a growing graph, one arc for each dependency the derivation establishes: an arc from chased to mouse, marking the internal argument; an arc from chased to cat, marking the external argument; and an arc from each occurrence of the to the noun it combines with, marking the specifier-of relation inside each nominal. The resulting object is the arc set $\{(chased, mouse), (chased, cat), (the1, cat), (the2, mouse)\}$, each arc labelled with the grammatical relation it encodes and each carrying its own direction. Unlike the nested sets K1 through K4 of Section 7, this graph does not need an external axiom to determine that chased selects mouse rather than the reverse, because the direction of selection is written into the arc itself; nor does Minimal Search need to guess whether a given node is a possible goal, since the search can simply follow arcs outward from a given probe and stop at the first node reached, exactly the operation Krivochen (2023) argues is well defined over digraphs and ill-defined over unordered sets. The price of this gain is ontological: a directed graph with labelled arcs is a considerably richer mathematical object than a bare set, and, as Zyman (2024) makes explicit, adopting it means abandoning the claim, central to the Minimalist Program's official self-description, that syntactic objects are nothing more than sets.

Taken together, this literature suggests that the choice between sets and graphs as the mathematical vehicle for Merge is not merely notational. A set-theoretic Merge must externalize order, directionality, and search through independent auxiliary axioms (the LCA, minimal search, accessibility conditions on the workspace); a graph-theoretic Merge can, in principle, build some of this information into the primitive structure-building operation itself, at the cost of committing to a richer mathematical object than Minimalism's official ontology has generally acknowledged.

Krivochen starts from the operational demands of Minimal Search and Zyman starts from the definitional demands of Merge itself; that these two independent lines of argument converge on the same graph-theoretic conclusion is, we think, more than a coincidence, and is treated further in Section 11.

4.8 Structural realism revisited: What survives theory change?

Returning to the apparatus of Section 2, we can now ask directly what a structural realist should say about the trajectory just traced. Votsis's (2011) continuity claim, that structural realists must account for which parts of a theory's structure survive replacement and which do not, without presupposing in advance that everything mathematical about a theory is equally invariant, gives us a concrete diagnostic to apply. C-command, we saw in Section 6, survives the shift from a representational definition over completed labelled trees to a derivational definition constituted by successive applications of Merge, with its empirical predictions essentially intact; this is exactly the kind of continuity-under-representational-change that structural realism, on Nefdt's (2021) application of it to generative linguistics, would predict for a genuine piece of grammatical reality rather than an artefact of a particular formalism. The deep/surface structure distinction, by contrast, did not survive the shift to Merge-based derivation at all (Section 5), and its disappearance cost the theory nothing empirically, the signature, on a structural realist reading, of formal overhead rather than structure that had latched onto anything real about language.

Other cases are less clear-cut, and it is here that the mathematical-structuralist point from Section 2, together with the model-based reading of grammars developed there, does real work. The unordered-ness of Merge (Section 7) and the free-algebra assumption challenged by multidominance and by Krivochen (2023) and Zyman (2024) (Section 9) are each, in their own way, candidates for either grade: unordered-ness is defended by an appeal to the Strong Minimalist Thesis and third-factor economy, which would make it a genuine discovery about the division of labour between narrow syntax and the interfaces if that argument holds up, but it could equally turn out, as deep structure did, to be an assumption forced by the choice of sets as the modelling vehicle rather than by anything about language itself. On Godfrey-Smith's (2006) resemblance criterion, the question to ask of any of these candidate models is not whether it is elegant but whether the respects in which it resembles the target, the actual dependencies speakers are sensitive to, are the respects doing the explanatory work, or merely artefacts of the vehicle chosen to carry them. Structural realism licenses optimism that some relational core will survive further reformulation of Merge; it does not, by itself, tell us in advance which parts of the current formalism constitute that core.

4.9 Grades of mathematical commitment

Drawing the preceding sections together, we can distinguish three grades of mathematical commitment that syntactic theorizing has, at different times, mistaken for one another. The first is descriptive adequacy: a mathematical object is adopted because it can state the generalizations linguists have independently discovered (constituency, c-command-sensitive dependencies, locality). Strings, trees, sets, algebras, and graphs have each, in their turn, achieved this grade for at least some empirical domain, which is why each could plausibly claim, in Rickart's (1995) deflationary sense of structure, to be capturing the same underlying relational pattern. The second grade is derivational economy: a mathematical object is preferred over a competitor because it derives, rather than stipulates, some property the field independently believes to be real, binarity derived from the LCA (Kayne 1994) rather than stipulated as a phrase-structure constraint is the clearest case, and it is a genuine advance because it reduces the number of independent parameters the theory must fix by hand. The third grade, and the one we think has generated the most confusion historically, is formal overhead: properties that a chosen mathematical object happens to have, or to require, purely in virtue of being that kind of object, with no independent linguistic motivation. The deep/surface distinction (Section 5) is the clearest historical casualty of this grade.

The critical upshot for current work is that the algebraic and graph-theoretic reformulations surveyed in Sections 8 and 9 should be evaluated by asking, for each proposed formal property, which of these three grades it belongs to, and this is not always transparent from the presentation of the proposals themselves. The Hopf-algebraic reformulations of Merge (Marcolli et al. 2023a, 2023b, 2025) are, at the time of writing, still substantially aspirational with respect to this distinction: it remains to be established which of the additional algebraic structure they import is forced by independently motivated linguistic considerations and which is formal overhead inherited from the physics contexts in which the relevant mathematics was first developed for other purposes. This is not a criticism unique to the algebraic programme: X-bar theory's stipulated two-bar-level template and the LCA's own assumption that c-command must be asymmetric were, at their respective moments of introduction, in exactly the same evidentiary position. It is, however, a caution the history recounted in this article suggests is worth making explicit rather than deferring to hindsight, as happened with deep structure. The convergence of Krivochen (2023) and Zyman (2024) on a graph-theoretic rather than set-theoretic Merge deserves particular attention on this score, precisely because the anomalies motivating it (set-theoretically ambiguous search configurations, discontinuous constituents) are not

obviously accommodable within the currently favoured formal object without additional derivational machinery, which is exactly the position deep structure occupied before Merge made its elimination possible.

V. CONCLUSION & RECOMMENDATIONS

5.1 Conclusion

We have traced the mathematical representation of syntactic structure from Saussure's associative sign-system, through Bloomfieldian and Harrisian distributionalism and Lambek's early categorial algebra, to the labelled graph of classical generative grammar, the unordered set of Merge-based Minimalism, and the graph-theoretic and algebraic reformulations current today: matrix syntax, Hopf-algebraic Merge, and multidominance. Before doing so, we argued that the terms “structure” and “structuralism” cover at least four distinct theses: linguistic structuralism as a school, mathematical structuralism as a position in the philosophy of mathematics, structural realism as a position in the philosophy of science, and the structuralist theory of science as a formal metatheoretic apparatus. This distinction clarifies rather than complicates the question of what is really being claimed when a syntactician says that a sentence “has” a certain structure. Grammars have always been models judged by their resemblance to speaker judgements, not literal transcripts of a formalism claiming psychological reality. Applied to the history itself, this framework suggests that the deep/surface structure distinction is the clearest casualty of formal overhead, that c-command's survival across representational change is the clearest case of a genuine relational discovery, and that the unordered-ness of Merge, the free-algebra assumption challenged by multidominance, and the additional structure imported by Hopf-algebraic and matrix-syntactic reformulations remain, at the time of writing, open cases whose ultimate status (derivational economy or formal overhead) the field is only beginning to be in a position to adjudicate.

5.2 Recommendations

Future research should systematically compare graph-theoretic, algebraic, and set-based representations of syntactic structure to determine which mathematical properties are necessary for explaining linguistic relations and which result from the formalism adopted. Particular attention should be given to the status of Merge, the explanatory role of c-command, and the implications of multidominance and linearization for syntactic theory. Such comparisons should prioritize descriptive adequacy and explanatory value while avoiding unnecessary formal complexity.

Declaration of Interest

The author declares that he does not have any known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Funding Declaration

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

REFERENCES

- Armstrong, D. M. (1978). *Nominalism and realism: Volume 1, Universals and scientific realism*. Cambridge University Press.
- Balzer, W., & Moulines, C. U. (Eds.). (1996). *Structuralist theory of science: Focal issues, new results*. De Gruyter. <https://doi.org/10.1515/9783110879421>
- Balzer, W., Sneed, J. D., & Moulines, C. U. (Eds.). (2000). *Structuralist knowledge representation: Paradigmatic examples*. Rodopi.
- Benacerraf, P. (1965). What numbers could not be. *The Philosophical Review*, 74(1), 47–73.
- Blevins, J. P. (1990). *Syntactic complexity: Evidence for discontinuity and multidomination* [Doctoral dissertation, University of Massachusetts Amherst]. Graduate Linguistics Student Association.
- Bloomfield, L. (1926). A set of postulates for the science of language. *Language*, 2(3), 153–164.
- Bloomfield, L. (1933). *Language*. Henry Holt.
- Carnie, A. (2008). *Constituent structure*. Oxford University Press.
- Chomsky, N. (1957). *Syntactic structures*. Mouton.
- Chomsky, N. (1965). *Aspects of the theory of syntax*. MIT Press.
- Chomsky, N. (1970). Remarks on nominalization. In R. A. Jacobs & P. S. Rosenbaum (Eds.), *Readings in English transformational grammar* (pp. 184–221). Ginn.
- Chomsky, N. (1981). *Lectures on government and binding*. Foris.
- Chomsky, N. (1986). *Barriers*. MIT Press.

- Chomsky, N. (1994). Bare phrase structure. In G. Webelhuth (Ed.), *Government and binding theory and the minimalist program* (pp. 383–439). Blackwell.
- Chomsky, N. (1995). *The minimalist program*. MIT Press.
- Chomsky, N. (2000). Minimalist inquiries: The framework. In R. Martin, D. Michaels, & J. Uriagereka (Eds.), *Step by step: Essays on minimalist syntax in honor of Howard Lasnik* (pp. 89–155). MIT Press.
- Chomsky, N. (2004). Beyond explanatory adequacy. In A. Belletti (Ed.), *Structures and beyond: The cartography of syntactic structures* (Vol. 3, pp. 104–131). Oxford University Press.
- Chomsky, N. (2013). Problems of projection. *Lingua*, 130, 33–49.
- Chomsky, N. (2015). Problems of projection: Extensions. In E. Di Domenico, C. Hamann, & S. Matteini (Eds.), *Structures, strategies and beyond: Studies in honour of Adriana Belletti* (pp. 3–16). John Benjamins.
- Chomsky, N., Seely, T. D., Berwick, R. C., Fong, S., Huybregts, M. A. C., Kitahara, H., McInnerney, A., & Sugimoto, Y. (2023). *Merge and the strong minimalist thesis*. Cambridge University Press. <https://doi.org/10.1017/9781009343244>
- Collins, C. (2002). Eliminating labels. In S. D. Epstein & T. D. Seely (Eds.), *Derivation and explanation in the minimalist program* (pp. 42–64). Blackwell.
- De Vries, M. (2009). On multidominance and linearization. *Biolinguistics*, 3(4), 344–403.
- Fox, D., & Pesetsky, D. (2005). Cyclic linearization of syntactic structure. *Theoretical Linguistics*, 31(1–2), 1–45.
- Godfrey-Smith, P. (2006). The strategy of model-based science. *Biology and Philosophy*, 21(5), 725–740.
- Harris, Z. S. (1951). *Methods in structural linguistics*. University of Chicago Press.
- Harris, Z. S., & Hiz, H. (1981). *Papers on syntax*. D. Reidel.
- Hellman, G., & Shapiro, S. (2019). *Mathematical structuralism*. Cambridge University Press.
- Jackendoff, R. (1977). *X-bar syntax: A study of phrase structure*. MIT Press.
- Kayne, R. S. (1994). *The antisymmetry of syntax*. MIT Press.
- Krivochen, D. G. (2023). The search for minimal search: A graph-theoretic approach. *Biolinguistics*, 17, Article e9793. <https://doi.org/10.5964/bioling.9793>
- Ladyman, J., & Ross, D. (2007). *Every thing must go: Metaphysics naturalized*. Oxford University Press.
- Lambek, J. (1958). The mathematics of sentence structure. *The American Mathematical Monthly*, 65(3), 154–170. <https://doi.org/10.1080/00029890.1958.11989160>
- Lieb, H.-H. (Ed.). (1992). *Prospects for a new structuralism*. John Benjamins.
- Marcolli, M., Berwick, R. C., & Chomsky, N. (2023a). *Old and new minimalism: A Hopf algebra comparison* [Preprint]. arXiv. <https://arxiv.org/abs/2306.10270>
- Marcolli, M., Berwick, R. C., & Chomsky, N. (2023b). *Syntax-semantics interface: An algebraic model* [Preprint]. arXiv. <https://arxiv.org/abs/2311.06189>
- Marcolli, M., Chomsky, N., & Berwick, R. C. (2025). *Mathematical structure of syntactic Merge: An algebraic model for generative linguistics*. MIT Press.
- Matthews, P. H. (2001). *A short history of structural linguistics*. Cambridge University Press.
- McCawley, J. D. (1982). Parentheticals and discontinuous constituent structure. *Linguistic Inquiry*, 13(1), 91–106.
- Meier, T. (2012). *A logical reconstruction of Leonard Bloomfield's linguistic theory* [Unpublished manuscript]. PhilSci-Archive. <https://philsci-archive.pitt.edu/9405/>
- Narita, H. (2014). *Endocentric structuring of projection-free syntax*. John Benjamins.
- Nefdt, R. M. (2016). *The foundations of linguistics: Mathematics, models, and structures* [Doctoral dissertation, University of St Andrews]. University of St Andrews Research Repository. <https://research-repository.st-andrews.ac.uk/handle/10023/9584>
- Nefdt, R. M. (2019). The philosophy of linguistics: Scientific underpinnings and methodological disputes. *Philosophy Compass*, 14(12), Article e12636. <https://doi.org/10.1111/phc3.12636>
- Nefdt, R. M. (2021). Structural realism and generative linguistics. *Synthese*, 199(1–2), 3711–3737. <https://doi.org/10.1007/s11229-020-02952-1>
- Newmeyer, F. J. (2023). Chomsky and the turn to syntax, including alternative approaches to syntax. In L. R. Waugh, M. Monville-Burston, & J. E. Joseph (Eds.), *The Cambridge history of linguistics* (pp. 549–576). Cambridge University Press. <https://doi.org/10.1017/9780511842788.030>
- Reinhart, T. (1976). *The syntactic domain of anaphora* [Doctoral dissertation, Massachusetts Institute of Technology].
- Resnik, M. D. (1997). *Mathematics as a science of patterns*. Oxford University Press.
- Rickart, C. E. (1995). *Structuralism and structures: A mathematical perspective*. World Scientific.
- Saussure, F. de. (1959). *Course in general linguistics* (W. Baskin, Trans.). Philosophical Library. (Original work published 1916)
- Shapiro, S. (1997). *Philosophy of mathematics: Structure and ontology*. Oxford University Press.



- Stabler, E. P. (1997). Derivational minimalism. In C. Retoré (Ed.), *Logical aspects of computational linguistics* (pp. 68–95). Springer.
- Votsis, I. (2011). Structural realism: Continuity and its limits. In P. Bokulich & A. Bokulich (Eds.), *Scientific structuralism* (pp. 105–117). Springer.
- Wells, R. S. (1947). Immediate constituents. *Language*, 23(2), 81–117.
- Worrall, J. (1989). Structural realism: The best of both worlds? *Dialectica*, 43(1–2), 99–124.
- Zyman, E. (2024). On the definition of Merge. *Syntax*, 27(4), 653–689. <https://doi.org/10.1111/synt.12287>.