Shift-grammars and foil creation

Diego Gabriel Krivochen & Beth Phillips

University of Reading

1. Introduction

The goal of this document is to provide the reader with a way to be able to create foil grammars for Artificial Grammar Learning (AGL) experimental paradigms using Lindenmayer systems in a systematic way such that:

- a. We can control for all surface statistical properties of strings¹, and
- b. The grammars generating the foils cannot be pruned to the grammar being probed for (a.k.a ‘target grammar’)

In other words: we want to be able to produce foil grammars which cannot be reduced by means of operations described elsewhere (Krivochen, 2017: 72; Uriagereka, 2015: 668) to the grammars being tested in an AGL paradigm of the kind used in Saddy (2009); Shirley (2014); Zormpa (2015); Phillips (2017); Vender et al. (forthcoming) among others. We will introduce the procedure without much formal discussion and keeping technicalities to a minimum: we aim at making this document useful in the design of experimental protocols rather than as a tool for the formal study of the generative capacity of Lindenmayer systems (but see Krivochen & Saddy, 2018 if formal detail is what you are after; also Krivochen & Phillips, 2018).

2. Procedure

We start by choosing an asymmetric grammar and determining its headedness: informally, and leaving aside technical aspects which are inconsequential to the present discussion, an asymmetric L-system is *right* or *left* headed depending on where the *strong* term (i.e., the term whose rewriting contains both itself and the other term) occurs in the rewriting function applying to the strong symbol. For instance:

Fib:

$$\begin{aligned} 0 &\rightarrow 1 \\ 1 &\rightarrow 0\ 1 \end{aligned}$$

This is *right-headed*: the ‘strong term’ is 1, since 1 rewrites as the other term and itself. In contrast, we can take a look at the bif grammar:

bif:

$$\begin{aligned} 0 &\rightarrow 1 \\ 1 &\rightarrow 1\ 0 \end{aligned}$$

This is *left-headed*.

We then proceed to develop the chosen grammar up to an arbitrary generation g_n in the traditional way: top-down, applying all rules simultaneously. But here’s the twist. In order to generate the following generation to g_n , g_{n+1} , we shift to a different procedure: because L-systems can be described by means of recurrence relations (such that any term g_n can be defined as a concatenation of previous terms, which may or may not be derivationally adjacent), a mechanistic procedure can build L-strings

¹ In the context of this paper and its siblings, a ‘string’ refers to an entire generation of the grammar under consideration.

by concatenating generations appropriately. In the case of Fib, we get the following recurrence relation:

$$Fib_n = Fib_{n-2} + Fib_{n-1}$$

Where ‘+’ is to be understood as linear concatenation: take generation $n-2$ and concatenate it to the left of generation $n-1$ to get target generation n . This follows from the definition of *headedness* above: 0 rewrites as 1, and 1 rewrites as g_{n-2} (i.e., 0) and g_{n-1} (i.e., 1).

If we develop the first three generations of Fib (note: the topmost 0 is not a generation, it is referred to as the ‘axiom’ of the grammar),

0
1
01
101

Generation 3 [101] can be expressed as the concatenation of generation 1 [1] to the left of generation 2 [01]. Headedness in L-systems is thus a *scale-free* property.

At this point, we can take a grammar G —say it’s *right-headed*, like Fib—, develop G up to some arbitrary generation n , and in order to derive $n+1$ *invert the concatenation order*. This inversion will be called a *shift*, and we will refer to the output of this inversion as a *shift grammar*. In other words: we get to generation n of G by right-concatenation, but in order to get to $n+1$ we shift to *left-concatenation*. Specifically, for Fib:

Before the shift: $Fib_n = Fib_{n-2} + Fib_{n-1}$

After the shift: $Fib_n = Fib_{n-1} + Fib_{n-2}$

Applied to our derivation above, this yields the following:

0
1
01
101 $\rightarrow n-2 + n-1$
10101 $\rightarrow n-1 + n-2$
10101101 $\rightarrow n-1 + n-2$
...

The last two lines of the derivation are the result of the shift.

3. Discussion: non-determinism and context-sensitivity

The rules of Fib are deterministic: a 0 always rewrites as a 1, and a 1 always rewrites as 01. Therefore, if we find a 1 followed by a 1, we know that the second one has to come from a 0, and the first one is the second symbol of a [01] constituent:

This will *always* be the case, since the bi-gram [00] is an illegitimate bi-gram in a Fib-derivation: the rules do *not* generate [00]. This means that we can, given a string corresponding to a *full* generation of Fib and the rules of the grammar, reliably reconstruct the previous generation of the grammar. The

possibility of having a reliable, unambiguous reconstruction procedure is essential in the choice of this particular kind of formalism for experimental purposes (as well as for purely formal research).

This is *not* the case when we apply the *head-shift*. Consider the relation between generations 4 and 5 in the shifted derivation above, repeated here:

$$\begin{aligned} 10101 &\rightarrow n-1 + n-2 \\ 1010101 &\rightarrow n-1 + n-2 \end{aligned}$$

We mark 1s preceded by 0s in red, and the rest in blue.

First of all, we need to note that the rewriting rule applying to 0 remains unchanged: 0 always rewrites as 1. But the shift has done interesting things to 1s: if we want to map g_4 to g_5 , we are forced to say that 1 can sometimes rewrite as [10] (marked in blue) and sometimes as [01] (marked in red). In practice, this means that we have moved into the space of *non-deterministic grammars*, for we need to amend the Fib-grammar as follows in order to account for the transition between these two generations:

$$\begin{aligned} 0 &\rightarrow 1 \\ 1 &\rightarrow 01 \\ 1 &\rightarrow 10 \end{aligned}$$

Now, we need to be able to specify the conditions under which each of the rules applies. Clearly we cannot simply look at single symbols as we have been doing with Fib, because there are two ways in which 1 can rewrite; this generates a situation in which there are two options:

- Assign probability values to each rule and let that rock
- Introduce an element of *context sensitivity* and keep the grammar under control

Option (b) amounts to enriching the rules in the following general way: ‘1 rewrites as [10] if it is followed by/preceded by ...x... and as [01] if it is followed by/preceded by ...y...’. Specifying a *shift-grammar*, then, would amount to making the contextual conditions fully explicit. It is crucial to note, however, that (b) only works if the grammar is *deterministic*. We will unpack this point in some detail. To give an example, we can take a closer look at our shift grammar. Descriptively, we can see that the 1s that rewrite as [01] are preceded by 0. Thinking about what they are *followed* by is not always useful, for the rightmost 1 is not followed by anything; we’d better make assumptions based on information we have. Likewise, the leftmost 1 is not *preceded* by anything, so that is not helpful. It *seems* like we can generalize on our first observation, and say the following:

$$1 \rightarrow 01 / 0_$$

Or, in words: ‘1 rewrites as [01] in the context of ‘preceded by 0’’. This is a *context-sensitive rule*, which means that a symbol rewrites as something depending not only on itself but also on the symbols that surround it. We can derive the following generation (g_6) of the *shift Fib* using the same concatenation procedure (which, remember, was $g_n = n-1 + n-2$) and see what happens:

$$\begin{aligned} &10101101 \\ &101011010101 \end{aligned}$$

We have marked the 0s that are relevant for the rewriting rule in green, and the 1s affected by that contextual condition in red. And we see that our prediction holds: 1s preceded by 0s indeed rewrite as [01] when we derive the following generation by the concatenation method. 0s always rewrite as 1, and that remains untouched. Can we say anything about 1s that rewrite as [10]? Yes, we can. In a sense, they can be thought of as the ‘elsewhere case’, which simplifies the grammar by encoding context-sensitivity and elsewhere-cases in a single rule:

$1 \rightarrow 01 / 0_; 10 \text{ elsewhere}$

This means ‘rewrite 1 as 01 when it is preceded by 0, and as 10 elsewhere’. And this indeed holds: we now identify the ‘elsewhere 1s’ in purple:

10101101

The prediction, then, is that the purple 1s will rewrite as (or ‘dominate’) [10]:

1010110110101

We have been consistent with the colour-coding: the input and the output of each rule have been coloured the same. Thus, the 0 that is referred to in the context-sensitive rule is green, and the 1 it generates is also green.

And the prediction holds, at least for the string we analyzed: we can now exhaustively analyse this generation of the shift-Fib grammar and reverse-engineer the previous generation in a fully deterministic way.

This is all well and good, but does it generalize? That is: can we say that *every Shift-grammar is a context-sensitive grammar*? In order to test this, we can apply the shift at a later generation, and see if the context-sensitivity we found holds. If it doesn’t, then we need to consider option (a), make the grammar non-deterministic (a.k.a. ‘stochastic’, or ‘probabilistic’). We applied the shift at generation 4, so we can ask: what happens if we jump a few generations?

Let’s now apply the shift to generation 7, to get generation 8:

$$\begin{aligned} g_7 &= \mathbf{101011010110110101101} \rightarrow n-2 + n-1 \\ g_8 &= \mathbf{10101101011011010110101101101} \rightarrow n-1 + n-2 \end{aligned}$$

We do not need to get too far into the string to find an *inconsistency*: consider the first 4-gram in g_7 : 1010, that is, two 1s followed by 0s. We would expect both of them to rewrite as the same bi-gram, given the fact that they appear in the same context, $_0$. However, as we can see in bold, this does not hold: the first 1 rewrites as [10] and the second one rewrites as [01]. In direct defiance of our context-sensitive rule, the following 1 (third one marked in bold), which is preceded by 1, *also rewrites as* [01]. To sum up: we have two 1s in distinct contexts –one preceded by 0 and the other preceded by 1– which rewrite *both* as [01]. Note that it would make no sense to say

$$\begin{aligned} 1 &\rightarrow 01 / 0_ \\ 1 &\rightarrow 01 / 1_ \end{aligned}$$

Because that leads to a contradiction: a 1 will always be preceded by either a 0 or a 1; that means that *every* 1 should rewrite as [01]...contrary to fact: there are 1s in shifted g_8 which rewrite as [10].

3.1 A frustrat(ed/ing) system

Is it possible to maintain the advantages of a deterministic system (context-sensitive as it may be) if we encode the context in the left-hand side of the rules? That is: imagine we managed to identify some n -gram length in shifted g_8 which yielded consistent rewritings. The reason is that there is more than a single possibility to define and identify constituency for local n -grams, and thus what gets rewritten can be more than one substring at a time: note that this is a problem in a rewriting system that is defined by the *simultaneous* application of all possible rules (and is thus ‘top-bottom’ rather than ‘left-to-right’, the latter being the usual state of affairs in Chomsky/Greibach-Normal grammars). Consider, for instance:

$$\begin{aligned} 01 &\rightarrow 101 \\ 10 &\rightarrow 0101 \end{aligned}$$

If we are presented with a substring generated by the grammar in (15), a deterministic parser cannot automatically decide what to rewrite, particularly because L-systems do not present the left-to-right directionality in rule application (a.k.a., the Traffic Convention) that sequential grammars do (including all versions of Phrase Structure Grammar): derivations in L-systems are strictly ‘top-down’, without any lateral orientation (and this is what allows all rules to apply simultaneously instead of sequentially). In L-systems there is no inherent hierarchy between the rules, that is, there is nothing that determines that rule X should apply *before* or *after* rule Y, contrarily to what happens in PSGs. In a system of the kind we analyze here, every 2-gram generates a conflict between both rules, which can –and thus, *must*– apply to the *same object at the same time*. Moreover, because we are dealing with a function-based automaton, there is no way of biasing the derivational process once it has started, consequently, what we get is effectively a situation in which two conflicting rules apply to the same input. A configuration in which there is a tension between mutually incompatible tendencies is referred to as a *dynamically frustrated system*, a term that comes from the study of ferromagnetic and anti-ferromagnetic interactions in electron lattices (Binder, 2008; Saddy, 2018). In our case, the *dynamical frustration* pertains to computational conditions over a local n -gram, and we can illustrate it as follows²:

...or this?

 Rewrite this...

The *frustration* illustrated above of course extends all throughout the string; this should serve as an argument for our claim that the left-hand side of L-rules cannot be tampered with if we aim to construct a deterministic L-system that does not require an external controller.

For completeness, let us take a look at the same procedure applied to bif. Recall, the bif grammar reverts the 1 rewriting rule, such that we are left with

0 → 1
 1 → 1 0

Now we develop this grammar:

0
 1
 10
 101
 10110
 10110101
 ...

The concatenation order in bif is the mirror image of Fib: generation n in bif results from concatenating $n-1 + n-2$, again *as per* the headedness properties of the grammar pointed out above. For the sake of comparison, we can apply the shift when deriving the generation after 101, which is what we did with the Fib grammar above. In that case, we have the following:

101 → $n-1 + n-2$
 10101 → $n-2 + n-1$

² This situation illustrates a kind of *computational frustration* that was not taken into consideration by Binder (2008), who focused only on deterministic Turing Machines. What we can show is that a computational frustration can arise under simpler conditions, and even outside the space defined by normal grammars.

$$\mathbf{10110101} \rightarrow n-2 + n-1$$

As above, the rewriting rule applying to 0 does not change: 0 always rewrites as 1. Again, if we looked at this fragment of the derivation we would be (mis)led into thinking that we have moved from deterministic context-free rules to context-sensitive rules. Let us look at this in some detail. What matters in the derivation of bif by concatenation is not what the relevant 1s are *preceded* by, but rather, what they are *followed by*: the mirror image of Fib. The relevant context-sensitive rule could then be:

$$1 \rightarrow 10 / _ 0$$

In words, ‘1 rewrites as 10 when it is *followed* by a 0’ (compare with the context-sensitive rule we formulated for Fib above, based on *precedence*). We can also jump ahead a bit and specify the elsewhere condition:

$$1 \rightarrow 10 / _ 0; 01 \text{ elsewhere}$$

And see if this works. As it happened with Fib, this apparent deterministic contextually determined rewritings break down when we attempt to generalize the shifts for arbitrary generations. The informal demonstration has the same form: we consider a shifted generation and identify symbols which, under the same contextual conditions, rewrite as different bigrams; if this happens, a context-sensitive system would need to include mutually contradictory rules. Again, we apply the shift to generate g_8 :

$$\begin{aligned} g_7 &= 10\mathbf{1101011011010110101} \rightarrow n-1 + n-2 \\ g_8 &= 101\mathbf{1010110110101101011011010110101} \rightarrow n-2 + n-1 \end{aligned}$$

We don’t need to go very far to find what we were looking for: the first [1] we highlighted occurs in the context 0_1 and rewrites as [10]; the second one occurs in the context 1_0, and *also* rewrites as [10]. Any hope for consistency in the formulation of context-sensitive rules must now be given up.

4. Some concluding remarks:

Let us summarise the argument in this paper. The result of applying a *shift* to the derivation of a grammar is a set of strings which maintain the same local transition probabilities as the vanilla grammar and the same grammaticality conditions: in the case of Fib and bif, these conditions are:

*111
*00

Crucially, the strings that result from the application of a shift *cannot be mapped to the vanilla grammar strings*, and this is precisely what makes these appropriate foils for experimental purposes *as per* the requirements stated in the Introduction. The longer the string, the better: quantitative properties of strings always manifest themselves more clearly at the limit (i.e., as the length of the string approaches infinity). For example:

Fib₆: 01**10110101101**
Shift-Fib₆: 10**10110110101**

Fib₈: 01**101101011011010110101101101101**
Shift-Fib₈: 10**10110101101101011010110110101101**

There are common substrings between the vanilla and the shifted versions given a specific generation (marked in **bold**), but it is important to point out that there is *no systematicity* to the distribution of these common n -grams between the grammars as generations unfold. In generation 8, for instance, surface properties are remarkably close, to the point that the strings only differ in a bi-gram. However,

hierarchical structure varies greatly, in non-trivial ways: Fib is a deterministic grammar and because of this it is possible to reconstruct previous generations by reverse-engineering the strings by means of a fully deterministic procedure: [01] always corresponds to [1] in the previous generation and [1] always corresponds to a [0] in the previous generation; we have seen that this does not hold for a *shift* string. It is *not possible* to reliably reconstruct structure in a *shift string* due to the fact that we would need to assume that 1 can rewrite as *both* [10] and [01] *in identical contexts*.

Importantly, common n -grams do *not* correspond to natural constituents in the grammars: in Fib, [10] is *not* a structural unit (because, remember, 1 rewrites as [01]; which means that the bi-gram [10] must come from a 0 and a 1). In Shift-Fib, however, [10] and [01] are both structural units, because they both are dominated by a single symbol (the linguists in the audience may refer to the definition of the relation ‘is-a’ in Immediate Constituent grammars). Having common superficial regularities yet differing structure-wise is an essential property that foils need to have, if AGL experimental protocols are probing for sensitivity to structure.

What implications does this hold for experimental application? If it is not possible to come up with deterministic context-sensitive rules which generalize for *any* generation of a shift grammar (in other words, the regularities that we may think we have found in g_n break down as we derive more generations), the only other option is to accept that determinism cannot be found in these grammars. And this has far-reaching consequences: most importantly, *it is not possible to reliably reconstruct hierarchical structure from a shift grammar*. This is an essential point for the AGL experimentalist; AGL tasks aim to assess whether the participant will be able to implicitly reconstruct the structure of an artificial grammar by exposure to grammatical strings alone. Previously, it has proven difficult to design L-system grammars which differ sufficiently in deep structure, but are not immediately distinct at surface level. When two deterministic grammars are increasingly similar at surface level, the probability that one can be locally assimilated to another may be higher. In other words: if the set of strings weakly generated by two grammars G and G' (say, G is the target grammar and G' is the foil) look similar in their surface properties (density of specific n -grams, ratio of 1s to 0s, local transition probabilities...), coming up with experimental conditions in which G and G' can be unambiguously told apart so that hypotheses about what is actually being probed for can be tested becomes more difficult. By keeping the surface statistical regularities of the grammar (its linear properties, low-order operations) constant, but messing up the core properties of structure reconstruction (i.e., hierarchy and constituency; higher-order operations) that Fib makes available (Krivochen, Phillips, & Saddy, 2018), it is possible to dissociate these two aspects in experimental designs. This is why we propose these shift grammars as excellent candidates for L-system foils in AGL tasks; consistent derivational rules cannot be developed generation by generation, and the previous generation simply cannot be reliably reconstructed, even if the participant was previously trained on the target grammar to the point of having implicitly learned some structure.

References:

- Binder, Philippe (2008) Frustration in Complexity. *Science* 320. 322-323.
- Krivochen, Diego Gabriel (2017) *Aspects of emergent cyclicity in language and computation*. PhD dissertation, University of Reading.
- Krivochen, Diego Gabriel & Beth Phillips (2018) A model for a Lindenmayer reconstruction algorithm. Ms. University of Reading.
- Krivochen, Diego Gabriel; Beth Phillips & Douglas Saddy (2018) Classifying points in Lindenmayer systems: transition probabilities and structure reconstruction. Ms. University of Reading.
- Krivochen, Diego Gabriel & Douglas Saddy (2018) Towards a classification of Lindenmayer systems. Ms. University of Reading. <https://arxiv.org/abs/1809.10542>

Lindner, Michael & Douglas Saddy (2017) *Lindenmeyer Systems Explorer* [Computer software] Reading, UK: Centre for Integrative Neurodynamics and Neuroscience, University of Reading.

Phillips, Beth (2017) *Symmetry of Hierarchical Artificial Grammars and Its Effects on Implicit Learning*. MSc dissertation. University of Reading.

Saddy, Douglas (2009) Perceiving and processing recursion in formal grammars. Paper Presented at *Recursion: Structural Complexity in Language and Cognition*. Conference at the University of Massachusetts (Amherst).

Saddy, Douglas (2018) Syntax and Uncertainty. In Angel Gallego & Roger Martin (eds.) *Language, Syntax, and the Natural Sciences*. Cambridge: CUP. 316-332.

Shirley, Elizabeth J. (2014) *Representing and Remembering Lindenmayer-Grammars*. PhD dissertation, University of Reading.

Uriagereka, Juan (2015) “Natural” grammars and natural language. In Beatriz Fernández and Pello Salaburu (eds.) *Ibon Sarasola, Gorazarre*. Bilbao: Editorial de la Universidad del País Vasco. 665-674.

Vender, Maria; Denis Delfitto; Douglas Saddy, Beth Phillips, & Diego Gabriel Krivochen (forthcoming) Implicit learning, bilingualism and dyslexia: Insights from a study assessing AGL with a modified Simon Task.

Zormpa, Eirini (2015) *Aspects of Bilingualism: Implicit learning and the interpretation of numerical quantifiers*. MSc dissertation, University of Reading.