Weighted Attribute Grammars: Reconciling Weighted and Attribute Grammars

Marcus Gerhold and Vadim Zaytsev
Formal Methods and Tools, University of Twente, The Netherlands

Weighted and probabilistic grammars offer an elegant way to rank, sample, or steer derivations. Attribute grammars offer an equally nice way to compute and propagate contextual information across a derivation tree. Either class is powerful on its own, yet there remains a useful class of problems in which quantitative choices depend directly on values computed during derivation. Here we introduce *Weighted Attribute Grammars* (WAG) as a single formalism that keeps both aspects explicit while keeping the specification short and conceptually clean.

A small but telling example is *binary exponential backoff*. After each collision, two communicating entities enlarge the contention window from which their next transmission attempt is chosen. The set of valid successful interaction traces is therefore infinite and context-sensitive: whether a pair such as $(4, 1)$ is admissible depends on how many earlier collisions occurred.

A plain weighted context-free grammar can encode bounded approximations of this behaviour, but the growth of the contention window then gets compiled away into a large family of productions. A plain attribute grammar can propagate the current window size, but then the quantitative side of the model has to be reconstructed in explicit semantic computations. Neither route is impossible. The difficulty is that the essential dependency of the model, since future quantitative choice depends on previously accumulated context, is not represented directly. This is precisely the dependency WAG make first-class.

A WAG combines a grammar skeleton with typed attributes, a weight domain, and production-local computations. Concretely, a WAG may be written as a tuple $\mathcal{G} = \langle N, T, P, s, \beta, A, \kappa, \pi, \Phi \rangle$, where the first four are nonterminals, terminals, production rules and the start symbol; β assigns each production rule either a literal weight or a weight expression over currently visible attributes, A is the set of typed attributes, κ and π assign inherited and synthesised attributes to nonterminals, and Φ provides the computations that update attribute values.

Each production rule has the schematic form $A\{\downarrow x, \uparrow y\} \xrightarrow{w} \alpha \ \langle \varphi \rangle$, where w is either a number or an expression over inherited or already available synthesised values, and φ updates attributes. The intended reading is simple: the choice of a production may depend on the current derivational context, and taking that production may in turn modify the context seen below or above it.

Context-free grammars embed in WAG as the case with neither attributes nor weights. Weighted or probabilistic context-free grammars embed as the case in which weights are present but do not depend on attributes. What is new is the interaction itself: contextual values may steer quantitative choice, and quantitative choice may participate in the evolution of contextual values.

Let us sketch a concise WAG for binary exponential backoff with four nonterminals S (start), P (protocol state), E (entity choice), and C (collision check):

$$\begin{aligned}
S &\xrightarrow{1} P\{[1,1]\}\\
P\{\downarrow v\} &\xrightarrow{1} E\{v,l\}\ E\{v,r\}\ C\{v,l,r\}\\
E\{\downarrow v, \uparrow r\} &\xrightarrow{v_i/|v|} \varepsilon\ \langle r := i\rangle\\
C\{\downarrow v, \downarrow x, \downarrow y\} &\xrightarrow{x=y} x, y;\ P\{\mathrm{dup}(v)\}\\
C\{\downarrow v, \downarrow x, \downarrow y\} &\xrightarrow{x\neq y} x, y.
\end{aligned}$$

Here v stores the current contention window, i ranges over its elements, each entity chooses an index with a weight determined by the current window, and the collision check either recurses with a doubled window or terminates successfully. The last two lines are guarded by conditions; operationally, one can read them as Boolean weights which enable productions exactly when the relation holds.

The modelling gain is not merely brevity. The growth of the contention window is stated once, where it logically belongs. The same artifact records both the admissible traces and the quantitative structure over them. Bounded encodings and auxiliary bookkeeping disappear from the specification and move back into the semantics, where they belong. This grammar is still recognisably close to context-free grammar notation: productions rewrite nonterminals into strings of terminals and nonterminals, and attributes flow in familiar inherited/synthesised directions. At the same time, the formalism can express dependencies that do not fit comfortably into ordinary weighted context-free grammars, because the derivation weight is no longer fixed statically by the production itself.

The same mechanism also scales beyond protocols, and applies whenever a grammar must "remember" some bounded or unbounded contextual quantity and use it to bias later derivations: resource-sensitive generation, adaptive test-data production, quantitative parsing policies, etc. WAG are not the only possible solution space, but they isolate a particularly direct and concise one.

Admitting interactions between attributes and weights reopens familiar metatheoretic questions in a more delicate form: *expressive placement* (which known grammar classes embed conservatively into WAG and how far attribute-sensitive weighting moves the formalism beyond weighted context-free grammars?); *consistency and termination* (in purely probabilistic grammars, even fixed production probabilities already require care to ensure that derivations induce a proper measure on terminal strings; in WAG, the relevant probabilities may vary with inherited attributes, so consistency conditions cannot simply be copied from the fixed-probability setting); *normalisation and parsing* (if one seeks CNF/GNF style normal forms, one must decide how much attribute structure survives such transformations and what parsing guarantees remain available once quantitative choice becomes attribute-sensitive; even when the underlying grammar skeleton remains context-free, the operational consequences of context-dependent weights are not obviously context-free any more).

To conclude, WAG offer a compact way to state quantitative context dependence inside the grammar. The binary exponential backoff case shows the intended payoff: quantitative and contextual parts remain in one place, and the spec stays close to what it describes. This makes WAG a good focus for further work on embeddings, consistency criteria, normal forms, and parsing algorithms.