



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Dynamic Programming Algorithms as Products of Weighted Logic Programs

Citation for published version:

Cohen, SB, Simmons, RJ & Smith, NA 2008, Dynamic Programming Algorithms as Products of Weighted Logic Programs. in *Proceedings of ICLP*. Springer, pp. 114-129.

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Proceedings of ICLP

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Dynamic Programming Algorithms as Products of Weighted Logic Programs

Shay B. Cohen, Robert J. Simmons, and Noah A. Smith

School of Computer Science
Carnegie Mellon University
{scohen, rjsimmon, nasmith}@cs.cmu.edu

Abstract. Weighted logic programming, a generalization of bottom-up logic programming, is a successful framework for specifying dynamic programming algorithms. In this setting, proofs correspond to the algorithm’s output space, such as a path through a graph or a grammatical derivation, and are given a weighted score, often interpreted as a probability, that depends on the score of the base axioms used in the proof. The desired output is a function over all possible proofs, such as a sum of scores or an optimal score. We describe the **PRODUCT** transformation, which can merge two weighted logic programs into a new one. The resulting program optimizes a product of proof scores from the original programs, constituting a scoring function known in machine learning as a “product of experts.” Through the addition of intuitive constraining side conditions, we show that several important dynamic programming algorithms can be derived by applying **PRODUCT** to weighted logic programs corresponding to *simpler* weighted logic programs.

1 Introduction

Weighted logic programming has found a number of applications in fields such as natural language processing, machine learning, and computational biology as a technique for declaratively specifying dynamic programming algorithms. Weighted logic programming is a generalization of bottom-up logic programming, with the numerical scores for proofs often interpreted as probabilities, implying that the weighted logic program implements *probabilistic* reasoning.

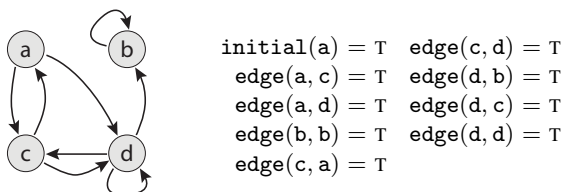
We describe a program transformation, **PRODUCT**, that is of special interest in weighted logic programming. **PRODUCT** transforms two weighted logic programs into a new one that implements probabilistic inference under an unnormalized probability distribution built as a product of the input programs’ distributions, known in machine learning as a “product of experts.” While this property has been exploited in a variety of ways in applications, there has not, to our knowledge, been a formal analysis or generalization in terms of the weighted logic programming representation.

The contribution of this paper is a general, intuitive, formal setting for dynamic programming algorithms that process two or more conceptually distinct structured inputs. Indeed, we show that many important dynamic programming algorithms can be derived using simpler “factor” programs and the **PRODUCT** transformation.

$$\text{reachable}(Q) \text{ :- } \text{initial}(Q). \quad (1)$$

$$\text{reachable}(Q) \text{ :- } \text{reachable}(P), \text{edge}(P, Q). \quad (2)$$

Fig. 1. A simple bottom-up logic program for graph reachability



$$\begin{array}{ll} \text{initial}(a) = T & \text{edge}(c, d) = T \\ \text{edge}(a, c) = T & \text{edge}(d, b) = T \\ \text{edge}(a, d) = T & \text{edge}(d, c) = T \\ \text{edge}(b, b) = T & \text{edge}(d, d) = T \\ \text{edge}(c, a) = T & \end{array}$$

Fig. 2. A directed graph and the corresponding initial database

The paper is organized as follows. In §2 we give an overview of weighted logic programming. In §3 we describe products of experts, a concept from machine learning that motivates our framework. In §4 we describe our framework and its connection to product of experts. In §5 we give derivations of several well-known algorithms using our framework.

2 Weighted Logic Programming

To motivate weighted logic programming, we begin with a logic program for single-source connectivity on a directed graph, shown in Fig. 1. In the usual bottom-up interpretation of this program, an initial database would describe the edge relation and one (or more) roots as axioms of the form $\text{initial}(a)$ for some a , and repeated forward inference would be applied on the two rules above to find the least database closed under those rules. However, in traditional logic programming this program can *only* be understood as a program calculating connectivity over a graph. Solving a different but structurally similar problem, such as a single-source shortest path, requires a rather different program to be written, and most solutions that have been presented require some form of non-deterministic committed choice [1,2].

Traditional logic programming is interpreted over Boolean values. A proof is a tree of valid inferences, and a valid proof is one where all of the leaves of the proof tree are axioms which are known to be true, and a true atomic proposition is one that has at least one valid proof. In weighted logic programming we generalize this notion to axioms, proofs, and atomic propositions “having values” rather than just “being true/valid.” A Boolean logic program takes the value of a proof to be the *conjunction* of the value of its axioms (so the proof has value “true” as long as all the propositions at the leaves are true-valued axioms), and takes the value of a proposition to be the *disjunction* of the values of its proofs (so an atomic proposition has value “true” if it has one true-valued proof). The single-source connectivity program would describe the graph in Fig. 2 by assigning T as the value of all the existing edges and the proposition $\text{initial}(a)$.

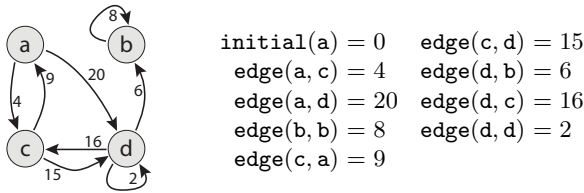


Fig. 3. A cost graph and the corresponding initial database

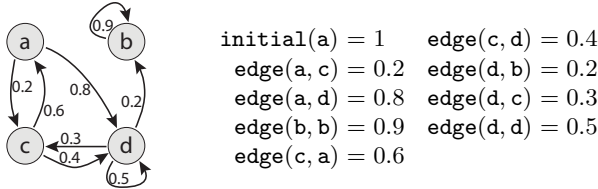


Fig. 4. A probabilistic graph and the corresponding initial database. With stopping probabilities made explicit, this would encode a Markov model.

2.1 Non-boolean Programs

With Weighted logic programming, the axioms and propositions can be understood as having non-Boolean values. In Fig. 3, axioms of the form $\text{edge}(X, Y)$ are given a value corresponding to the cost along the edge in the graph, and the axiom $\text{initial}(a)$ is given the value 0. If we take the value or “score” of a proof to be the *sum* of the values of its axioms, and then take the value of a proposition to be the *minimum* score over all possible proofs, then the program from Fig. 1 describes *single-source shortest path*. We replace the operators :- (disjunction) and , (conjunction) with $\text{min} =$ and $+$, respectively, and interpret the program over the non-negative numbers. With a specific execution strategy, the result is Dijkstra’s single-source shortest path algorithm.

Whereas Fig. 3 describes a cost graph, in Fig. 4 weights on edges are to be interpreted as *probabilities*, so that the graph can be seen as a Markov model or probabilistic finite-state network over which random walks are well-defined.¹ If we replace :- (disjunction) and , (conjunction) with $\text{max} =$ and \times , then the value of $\text{reachable}(X)$ for any X is the probability of the most likely path from a to X . For instance, $\text{reachable}(a)$ ends up with the value 1, and $\text{reachable}(b)$ ends up with value 0.16, corresponding to the path from $a \rightarrow d \rightarrow b$, whose weight is $\text{initial}(a) \times \text{edge}(a, d) \times \text{edge}(d, b)$.

If we keep the initial database from Fig. 4 but change our operators from $\text{max} =$ and \times to $+=$ and \times , the result is a program for *summing* over the probabilities of all distinct paths that start in a and lead to X , for each vertex X . This quantity is known as the “path sum” [3]. The path sum for b , for instance, is 10—this is not a probability, but rather a sum of probabilities of many paths, some of which are prefixes of each other.

¹ For each vertex, the out-going edges’ weights must be non-negative and sum to a value less than or equal to one. Remaining probability mass is assumed to go to a “stopping” event, as happens with probability 0.1 in vertex b in Fig. 4.

$$\text{reachable}(\mathbf{Q}) \oplus = \text{initial}(\mathbf{Q}). \quad (3)$$

$$\text{reachable}(\mathbf{Q}) \oplus = \text{reachable}(\mathbf{P}) \otimes \text{edge}(\mathbf{P}, \mathbf{Q}). \quad (4)$$

Fig. 5. The logic program from Fig. 1, rewritten to emphasize that it is generalized to an arbitrary semiring

These three related weighted logic programs are useful generalizations of the reachability logic program. Fig. 5 gives a generic representation of all four algorithms in the Dyna language [4]. The key difference among them is the *semiring* in which we interpret the weights.² Reachability uses the Boolean semiring $\langle \{\mathbf{T}, \mathbf{F}\}, \vee, \wedge, \mathbf{F}, \mathbf{T} \rangle$, single-source shortest path uses $\langle \mathbb{R}_{\geq 0} \cup \{\infty\}, \min, +, \infty, 0 \rangle$, the most-probable path variant uses $\langle [0, 1], \max, \times, 0, 1 \rangle$, and the probabilistic path-sum variant uses $\langle \mathbb{R}_{\geq 0} \cup \{\infty\}, +, \times, 0, 1 \rangle$. The formalism we describe here requires semirings that are closed under finite products and infinite sums—in our running example, this corresponds to the idea that there may be infinite paths through a graph, all with finite length.

Weighted logic programming arose in the computational linguistics community [5] after it was argued by Shieber, Schabes, and Pereira [6] and Sikkal [7] that many parsing algorithms for non-deterministic grammars could be represented as deductive logic programs, and McAllester [8] showed that this representation facilitates reasoning about asymptotic complexity. Other developments include a connection between weighted logic programs and hypergraphs [9], optimal A* search for maximizing programs [10], semiring-general agenda-based implementations [11], improved k -best algorithms [12], and program transformations to improve efficiency [13].

2.2 Formal Definition

A weighted logic program is a set of “Horn equations” [13] describing a set of declarative, usually recursive equations over an abstract semiring:

$$\text{consequent}(\mathbf{U}) \oplus = \text{antecedent}_1(\mathbf{W}_1) \otimes \dots \otimes \text{antecedent}_n(\mathbf{W}_n).$$

Here \mathbf{U} and the \mathbf{W}_i are sequences of variables X_1, \dots, X_k . If $\mathbf{U} \subseteq \bigcup_{i=1}^n \mathbf{W}_i$ for every rule, then the program is *range-restricted* or *fully grounded*.

A weighted logic program is specified on an arbitrary semiring, and can be interpreted in any semiring $\langle \mathbb{K}, \oplus, \otimes, \mathbf{0}, \mathbf{1} \rangle$, as previously described.

Typically the proof and value of a *specific* theorem are desired. We assume that this theorem is called `goal` and takes zero arguments. A computationally uninteresting but perfectly intuitive way to present a weighted logic program is

$$\text{goal} \oplus = \text{axiom}_1(\mathbf{W}_1) \otimes \dots \otimes \text{axiom}_n(\mathbf{W}_n).$$

The value of the proposition/theorem `goal` is a semiring-sum over all of its proofs, starting from the axioms, where the value of any single proof is the semiring-product of

² An algebraic semiring consists of five elements $\langle \mathbb{K}, \oplus, \otimes, \mathbf{0}, \mathbf{1} \rangle$, where \mathbb{K} is a domain closed under \oplus and \otimes , \oplus is a binary, associative, commutative operator, \otimes is a binary, associative operator that distributes over \oplus , $\mathbf{0} \in \mathbb{K}$ is the \oplus -identity, and $\mathbf{1} \in \mathbb{K}$ is the \otimes -identity.

the axioms involved. This is effectively encoded using the inference rules as a sum of products of sums of products of ... sums of products, exploiting distributivity and shared substructure for efficiency. Dynamic programming algorithms, useful for problems with a large degree of shared substructure, are often encoded as weighted logic programs.

In many practical applications, as in our reachability example in Section 2.1, values are interpreted as probabilities to be maximized or summed or costs to be minimized.

3 Products of Experts

In machine learning, probability models learned from example data are often used to make predictions. For example, to predict the value of a random variable Y (ranging over values denoted y in a domain denoted \mathcal{Y} ; here, \mathcal{Y} corresponds to the set of proofs) given that random variable X has an observed value $x \in \mathcal{X}$ (here, \mathcal{X} ranges over initial databases, i.e., sets of axioms), the Bayes decision rule predicts:

$$\hat{y} = \operatorname{argmax}_{y \in \mathcal{Y}} p(Y = y \mid X = x) = \operatorname{argmax}_{y \in \mathcal{Y}} \frac{p(Y = y, X = x)}{p(X = x)} \quad (5)$$

In other words, the prediction \hat{y} should be the element of \mathcal{Y} that maximizes $p(Y = y \mid X = x)$, the likelihood of the event $Y = y$ given that the event $X = x$ has happened. By the definition of conditional probability, this quantity is equivalent to the ratio of the *joint* probability that $X = x \wedge Y = y$ to the *marginal* probability that $X = x$. Dynamic programming algorithms are available for solving many of these maximization problems, such as when Y ranges over paths through a graph or grammar derivations.

Of recent interest are probability models p that take a *factored form*, for example:

$$p(X = x, Y = y) \propto p_1(X = x, Y = y) \times \dots \times p_n(X = x, Y = y) \quad (6)$$

where \propto signifies “proportional to” and suppresses the means by which the probability distribution is re-normalized to sum to one. This kind of model is called a *product of experts* [14]. Intuitively, the probability of an event under p can only be large if “all the experts concur,” i.e. if the probability is large under each of the p_i . Any single expert can make an event arbitrarily unlikely (even impossible) by giving it very low probability.

The attraction of such probability distributions is that they modularize complex systems [15,16]. They can also offer computational advantages when solving Eq. 5 [17]. Further, the expert factors can often be trained (i.e., estimated from data) separately, speeding up expensive but powerful machine learning methods [18,19].

This idea is still useful even when not dealing with probabilities. Suppose each expert p_i is a function $\mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$ that returns 1 if and only if the arguments x and y satisfy some constraints; it implements a relation. Then the “product” relation is just the intersection of all pairs $\langle x, y \rangle$ for which all the expert factors’ relations hold.

To the best of our knowledge, there has been no attempt to formalize the following intuitive idea about products of experts: algorithms for summing and maximizing mutually-constrained pairs of product-proof values should resemble the individual algorithms for each of the two separate proofs’ values. Our formalization is intended to aid in algorithm

$$\text{reachable}_1(Q_1) \oplus= \text{initial}_1(Q_1). \quad (7)$$

$$\text{reachable}_1(Q_1) \oplus= \text{reachable}_1(P_1) \otimes \text{edge}_1(P_1, Q_1). \quad (8)$$

$$\text{reachable}_2(Q_2) \oplus= \text{initial}_2(Q_2). \quad (9)$$

$$\text{reachable}_2(Q_2) \oplus= \text{reachable}_2(P_2) \otimes \text{edge}_2(P_2, Q_2). \quad (10)$$

Fig. 6. Two identical “experts” for generalized graph reachability, duplicates of the program in Fig. 5

$$\text{reachable}_{1 \circ 2}(Q_1, Q_2) \oplus= \text{initial}_1(Q_1) \otimes \text{initial}_2(Q_2). \quad (11)$$

$$\text{reachable}_{1 \circ 2}(Q_1, Q_2) \oplus= \text{reachable}_2(P_2) \otimes \text{edge}_2(P_2, Q_2) \otimes \text{initial}_1(Q_1). \quad (12)$$

$$\text{reachable}_{1 \circ 2}(Q_1, Q_2) \oplus= \text{reachable}_1(P_1) \otimes \text{edge}_1(P_1, Q_1) \otimes \text{initial}_2(Q_2). \quad (13)$$

$$\text{reachable}_{1 \circ 2}(Q_1, Q_2) \oplus= \text{reachable}_{1 \circ 2}(P_1, P_2) \otimes \text{edge}_1(P_1, Q_1) \otimes \text{edge}_2(P_2, Q_2). \quad (14)$$

Fig. 7. Four rules that, in addition to the rules in Fig. 6, define the product of experts of reachable_1 and reachable_2

development as new kinds of random variables are coupled, with a key practical advantage: the expert factors are known because they fundamentally *underlie* the main algorithm. Indeed, we call our algorithms “products” because they are derived from “factors.”

4 Products of Weighted Logic Programs

In this section, we will motivate products of weighted logic programs in the context of the running example of generalized graph reachability. We will then define the **PRODUCT** transformation precisely and describe the process of specifying new algorithms as constrained versions of product programs.

Fig. 6 defines two “experts,” copies of the graph reachability program from Fig. 5. We are interested in a new predicate $\text{reachable}_{1 \circ 2}(Q_1, Q_2)$, which for any particular X and Y should be equal to the product of $\text{reachable}_1(X)$ and $\text{reachable}_2(Y)$. We could define the predicate by adding the following rule to the program in Fig. 6:

$$\text{reachable}_{1 \circ 2}(Q_1, Q_2) \oplus= \text{reachable}_1(Q_1) \otimes \text{reachable}_2(Q_2).$$

This program is a bit simplistic, however; it merely describes calculating the “experts” independently and then combining them at the end. The key to the **PRODUCT** transformation is that the predicate of $\text{reachable}_{1 \circ 2}$ can alternatively be calculated by adding the following four rules to Fig. 6:

$$\begin{aligned} \text{reachable}_{1 \circ 2}(Q_1, Q_2) &\oplus= \text{initial}_1(Q_1) \otimes \text{initial}_2(Q_2). \\ \text{reachable}_{1 \circ 2}(Q_1, Q_2) &\oplus= \text{initial}_1(Q_1) \otimes \text{reachable}_2(P_2) \otimes \text{edge}_2(P_2, Q_2). \\ \text{reachable}_{1 \circ 2}(Q_1, Q_2) &\oplus= \text{reachable}_1(P_1) \otimes \text{edge}_1(P_1, Q_1) \otimes \text{initial}_2(Q_2). \\ \text{reachable}_{1 \circ 2}(Q_1, Q_2) &\oplus= \text{reachable}_1(P_1) \otimes \text{edge}_1(P_1, Q_1) \otimes \\ &\quad \text{reachable}_2(P_2) \otimes \text{edge}_2(P_2, Q_2). \end{aligned}$$

Then, because $\text{reachable}_{1 \circ 2}(Q_1, Q_2)$ was defined above to be the product of $\text{reachable}_1(Q_1)$ and $\text{reachable}_2(Q_2)$, it should be clear that the last rule can be

rewritten to obtain the factored program in Fig. 7. This program computes over pairs of paths in two graphs.

4.1 The PRODUCT Transformation

The PRODUCT program transformation is shown in Fig. 8. For each desired product of experts, where one “expert,” the predicate p , is defined by n rules and the other expert q by m rules, the transformation defines the product of experts for $p \circ q$ with $n \times m$ new rules, the cross product of inference rules from the first and second experts. The value of a coupled proposition $p \circ q$ in \mathcal{P}' will be equal to the semiring product of p 's value and q 's value in \mathcal{P} (or, equivalently, in \mathcal{P}').

Input: A logic program \mathcal{P} and a set \mathcal{S} of pairs of predicates (p, q) .

Output: A program \mathcal{P}' that extends \mathcal{P} , additionally computing the product predicate $p \circ q$ for every pair $(p, q) \in \mathcal{S}$ in the input.

```

1:  $\mathcal{P}' \leftarrow \mathcal{P}$ 
2: for all pairs  $(p, q)$  in  $\mathcal{S}$  do
3:   for all rules in  $\mathcal{P}$ , of the form  $p(\mathbf{W}) \oplus= A_1 \otimes \dots \otimes A_n$  do
4:     for all rules in  $\mathcal{P}$ , of the form  $q(\mathbf{X}) \oplus= B_1 \otimes \dots \otimes B_m$  do
5:       let  $r \leftarrow [p \circ q(\mathbf{W}, \mathbf{X}) \oplus= A_1 \otimes \dots \otimes A_n \otimes B_1 \otimes \dots \otimes B_m]$ 
6:       for all pairs of antecedents in  $r$   $(\mathbf{s}(\mathbf{Y}), \mathbf{t}(\mathbf{Z}))$  such that  $(\mathbf{s}, \mathbf{t}) \in \mathcal{S}$  do
7:         remove the antecedents  $\mathbf{s}(\mathbf{Y})$  and  $\mathbf{t}(\mathbf{Z})$  from  $r$ 
8:         insert the antecedent  $\mathbf{s} \circ \mathbf{t}(\mathbf{Y}, \mathbf{Z})$  to  $r$ 
9:       end for
10:      add  $r$  to  $\mathcal{P}'$ 
11:     end for
12:   end for
13: end for
14: return  $\mathcal{P}'$ 

```

Fig. 8. This figure describes PRODUCT, a non-deterministic program transformation that adds new rules to WLP \mathcal{P} that compute the product of experts of predicates from the original program. We implicitly rename variables to avoid conflicts between rules.

Note that lines 6–8 are non-deterministic under certain circumstances, because if the antecedent of the combined program is $\mathbf{a}(\mathbf{X}) \otimes \mathbf{a}(\mathbf{Y}) \otimes \mathbf{b}(\mathbf{Z})$ and the algorithm is computing the product of \mathbf{a} and \mathbf{b} , then the resulting antecedent could be either $\mathbf{a} \circ \mathbf{b}(\mathbf{X}, \mathbf{Z}) \otimes \mathbf{a}(\mathbf{Y})$ or $\mathbf{a} \circ \mathbf{b}(\mathbf{Y}, \mathbf{Z}) \otimes \mathbf{a}(\mathbf{X})$. Our procedure arbitrarily selects one of the possibilities.

4.2 Constraining the Product of Experts

Any program \mathcal{P}' that comes out after applying PRODUCT on \mathcal{P} computes the product of experts of p and q (where $(p, q) \in \mathcal{S}$). More specifically, any ground instances of $p(\mathbf{X})$ and $q(\mathbf{Y})$ have the same value in \mathcal{P} and \mathcal{P}' , and the value of $p \circ q(\mathbf{X}, \mathbf{Y})$ in \mathcal{P}' is $p(\mathbf{X}) \otimes q(\mathbf{Y})$.³ However, the same program could have been implemented more straightforwardly by merely introducing a new inference rule for the goal.

³ The full proof can be found in [20].

$$\text{reachable}_{e_{1o2}}(Q_1, Q_2) \oplus = \text{initial}_1(Q_1) \otimes \text{initial}_2(Q_2). \quad (15)$$

$$\text{reachable}_{e_{1o2}}(Q_1, Q_2) \oplus = \text{reachable}_{e_{1o2}}(P_1, P_2) \otimes \text{edge}_1(P_1, Q_1) \otimes \text{edge}_2(P_2, Q_2). \quad (16)$$

Fig. 9. By removing all but these two rules from the product of experts in Fig. 7, we require both paths to have the same number of steps

$$\text{reachable}_{e_{1o2}}(Q) \oplus = \text{initial}_1(Q) \otimes \text{initial}_2(Q). \quad (17)$$

$$\text{reachable}_{e_{1o2}}(Q) \oplus = \text{reachable}_{e_{1o2}}(P) \otimes \text{edge}_1(P, Q) \otimes \text{edge}_2(P, Q). \quad (18)$$

Fig. 10. By further constraining the program in Fig. 9 to require that the $Q_1 = Q_2$ at all points, we require both paths to be identical

Yet, the output of the PRODUCT transformation is a starting point for describing dynamic programming algorithms that do two similar actions—traversing a graph, scanning a string, parsing a sentence—at the same time and in a coordinated fashion. Exactly what “coordinated fashion” means depends on the problem, and answering that question determines how the problem is constrained.

If we return to the running example of generalized graph reachability, the program as written has eight rules, four from Fig. 6 and four from Fig. 7. Two examples of constrained product programs are given in Fig. 9 and Fig. 10. In the first example in Fig. 9, the only change is that all but two rules have been removed from the program in Fig. 7. Whereas in the original product program $\text{reachable}_{e_{1o2}}(Q_1, Q_2)$ corresponded to the product of the weight of the “best” path from the initial state or states of graph 1 to Q_1 and the weight of the “best” path from the initial state or states of graph 2 to Q_2 , the new program computes the best paths from the two origins to the two destinations with the additional requirement that the paths be the *same length*—the rules that were deleted allowed for the possibility of a prefix on one path or the other.

If our intent is for the two paths to not only have the same length but to visit vertices in the same sequence, then we can further constrain the program to only define $\text{reachable}_{e_{1o2}}(Q_1, Q_2)$ where $Q_1 = Q_2$, at which point it might as well be written $\text{reachable}_{e_{1o2}}(Q)$. This is what is done in Fig. 10.

The choice of paired predicates \mathcal{S} is important for the final WLP which PRODUCT returns and it also limits the way we can add constraints to derive a new WLP. Automatically deriving \mathcal{S} from data in a machine learning setting is an open question for future research. When PRODUCT is applied on two copies of the same WLP (concatenated together to a single program), a natural schema for selecting paired predicates arises, in which we pair a predicate from one program with the same predicate from the other program. This natural pairing leads to the derivation of several useful, known algorithms, to which we turn in Section 5.

5 Examples

In this section, we describe three classes of algorithms that can be understood as constrained products of simpler weighted logic programs.

$$\text{dist}(P) \oplus = \text{start}(P). \quad (19)$$

$$\text{dist}(P) \oplus = \text{dist}(P) \otimes \text{staycost}. \quad (20)$$

$$\text{dist}(P + 1) \oplus = \text{dist}(P) \otimes s(C, P). \quad (21)$$

Fig. 11. A program for scanning over a string

$$\text{dist}_{1 \circ 2}(P_1, P_2) \oplus = \text{start}_1(P_1) \otimes \text{start}_2(P_2). \quad (22)$$

$$\text{dist}_{1 \circ 2}(P_1, P_2 + 1) \oplus = \text{start}_1(P_1) \otimes \text{dist}_2(P_2) \otimes s(C_2, P_2). \quad (23)$$

$$\text{dist}_{1 \circ 2}(P_1 + 1, P_2) \oplus = \text{dist}_1(P_1) \otimes s(C_1, P_1) \otimes \text{start}_2(P_2). \quad (24)$$

$$\text{dist}_{1 \circ 2}(P_1, P_2 + 1) \oplus = \text{dist}_{1 \circ 2}(P_1, P_2) \otimes s_2(C_2, P_2) \otimes \text{staycost}_1. \quad (25)$$

$$\text{dist}_{1 \circ 2}(P_1 + 1, P_2) \oplus = \text{dist}_{1 \circ 2}(P_1, P_2) \otimes s_1(C_1, P_1) \otimes \text{staycost}_2. \quad (26)$$

$$\text{dist}_{1 \circ 2}(P_1 + 1, P_2 + 1) \oplus = \text{dist}_{1 \circ 2}(P_1, P_2) \otimes s_2(C_1, P_1) \otimes s_2(C_2, P_2) \boxed{\text{if } C_1 = C_2}. \quad (27)$$

Fig. 12. Edit distance derived from the PRODUCT transformation on two copies of Fig. 11, with a side condition (boxed)

5.1 Edit Distance

Edit distances [21] are important measures of the difference between two strings, and they underlie many algorithms in computational biology and computational linguistics. The DNA fragment “ACTAGCACTTAG” can be encoded as a set of axioms $s(a, 1)$, $s(c, 2)$, $s(t, 3)$, \dots , $s(g, 12)$, and we can describe a specification of the dynamic program for edit distance by using the product of a trivial “cursor” program that scans over a string, described in Fig. 11.

We generally interpret Fig. 11 over the the “cost minimization” semiring, replacing $\oplus =$ with $\min =$ and \otimes with $+$. The value of all the axioms of the form $\text{start}(P)$ (giving the starting position) or $s(C, P)$ (indicating the position of a character) is 0, but the value of staycost is some finite, nonzero value representing the penalty if the cursor stays in one place. Note that under this semiring interpretation, the rule (20) will never be used; the value of staycost only becomes relevant when the PRODUCT of the scanning program with itself is determined.

The output of PRODUCT on two copies of the scanning program is shown in Fig. 12, though three rules are removed.⁴ One important change is made to clause 27, the addition of the *side condition* $C_1 = C_2$, which requires that when both P_1 and P_2 advance, the character at position P_1 in string 1 and the character at position P_2 in string two must be identical. This captures the essential requirement of the edit distance calculation: changing a symbol in one string to an identical symbol in the other string incurs no “edit” and no cost. It is worth noting that the clause 27 could have equivalently been written by unifying C_1 with C_2 :

$$\text{dist}_{1 \circ 2}(P_1 + 1, P_2 + 1) \oplus = \text{dist}_{1 \circ 2}(P_1, P_2) \otimes s_2(C, P_1) \otimes s_2(C, P_2).$$

⁴ The three removed rules are the two different combinations of clause 19 and clause 20 in Fig. 11, as well as the combination of clause 20 with itself. These three rules are redundant if we are computing in the minimum-cost semiring.

The first and last clauses in the edit distance program (22 and 27) are the essential ones; the other five clauses essentially describe extensions to edit distance that can be added, turned off, or modified to obtain different edit distances. Clause 23 allows a penalty-free prefix to be added to the second string, for instance matching string 1, “BOVIK” against string 2, “HARRY BOVIK,” and clause 24 allows a penalty-free prefix to be deleted from the first string. Clause 25 describes that insertions can be made with cost staycost_1 , for instance in matching string one “H. BOVIK” against “H. Q. BOVIK”, and clause 26 describes deletions from the first string to reach the second.

5.2 Finite-State Algorithms

Specifications of weighted finite-state automata (WFSAs) and transducers (WFSTs) are superficially similar to the reachability problem of Sec. 1, but with edge relations ($\text{edge}(P, Q)$) augmented by symbols (WFSAs: $\text{arc}(P, Q, A)$) or pairs of input-output symbols (WFSTs: $\text{arc}(P, Q, A, B)$). Weighted finite-state machines are widely used in speech and language processing [22].

Weighted Finite State Automata. Fig. 13 describes an algorithm for recognizing paths in a weighted finite state automaton. (With the appropriate semirings, it finds the most probable path or the path-sum.) If the **PRODUCT** of that algorithm with itself is taken, we can follow similar steps in Sec. 4.2 and add a constraint to clause 33 that requires the two paths’ symbols to be identical, we get the recognizer for the (*weighted*) *intersection* of the two WFSAs (itself a WFA). Weighted intersection generalizes intersection, and can be used, for example, to determine whether a specific string (itself an FSA) is in the regular language of the FSA and, in the probabilistic case, its associated probability.

Weighted Finite-State Transducers. Suppose we take the **PRODUCT** transformation of the WFST recognition algorithm (not shown, but similar to Fig. 13 but using $\text{arc}(P, Q, A, B)$ as arc axioms) with itself and constrain the result by removing all but the three interesting rules (as before) and requiring that B_1 (the “output” along the first edge) always be equal to A_2 (the “input” along the second edge). The result is shown

$$\text{goal} \oplus= \text{path}(Q) \otimes \text{final}(Q). \quad (28)$$

$$\text{path}(Q) \oplus= \text{initial}(Q). \quad (29)$$

$$\text{path}(Q) \oplus= \text{path}(P) \otimes \text{arc}(P, Q, A). \quad (30)$$

Fig. 13. The weighted logic program describing (weighted) recognition in a probabilistic finite state automaton

$$\text{goal}_{1 \otimes 2} \oplus= \text{path}_{1 \otimes 2}(Q_1, Q_1) \otimes \text{final}_1(Q_2) \otimes \text{final}_2(Q_2). \quad (31)$$

$$\text{path}_{1 \otimes 2}(Q_1, Q_2) \oplus= \text{initial}_1(Q_1) \otimes \text{initial}_2(Q_2). \quad (32)$$

$$\text{path}_{1 \otimes 2}(Q_1, Q_2) \oplus= \text{path}_{1 \otimes 2}(P_1, P_2) \otimes \text{arc}_1(P_1, Q_1, A_1) \otimes \text{arc}_2(P_2, Q_2, A_2) \boxed{\text{if } A_1 = A_2.} \quad (33)$$

Fig. 14. The weighted logic program describing (weighted) recognition by an intersection of two finite state automata, derived from Fig. 13 in the manner of Fig. 9

$$\text{goal}_{1\circ 2} \oplus = \text{path}_{1\circ 2}(Q_1, Q_1) \otimes \text{final}_1(Q_2) \otimes \text{final}_2(Q_2). \quad (34)$$

$$\text{path}_{1\circ 2}(Q_1, Q_2) \oplus = \text{initial}_1(Q_1) \otimes \text{initial}_2(Q_2). \quad (35)$$

$$\text{path}_{1\circ 2}(Q_1, Q_2) \oplus = \text{path}_{1\circ 2}(P_1, P_2) \otimes \text{arc}_1(P_1, Q_1, A_1, B_1) \otimes \text{arc}_2(P_2, Q_2, A_2, B_2) \boxed{\text{if } B_1 = A_2.}$$

Fig. 15. The weighted logic program describing a composition of two finite state transducers, derived from Fig. 13 in the manner of Fig. 9 and Fig. 14

in Fig. 15; this is the recognition algorithm for the WFST resulting from *composition* of two WFSTs. Composition permits small, understandable components to be cascaded and optionally compiled, forming complex but efficient models over strings.

5.3 Context-Free Parsing

Parsing natural languages is a difficult, central problem in computational linguistics [23]. Consider the sentence “Shay saw Robert with binoculars.” One analysis (the most likely in the real world) is that Shay had the binoculars and saw Robert through them. Another is that *Robert* had the binoculars, and Shay saw the binocular-endowed Robert. Fig. 16 shows syntactic parses into noun phrases (**NP**), verb phrases (**VP**), etc., corresponding to these two meanings. It also shows part of a context-free grammar describing English sentences in Chomsky normal form [24],⁵ and an encoding of the grammar and string using axioms. A proof corresponds to a CF derivation of the string.

In [6], the authors show that parsing with CFGs (and other grammars) can be formalized as a logic program, and in [5] this framework is extended to the weighted case. If weights are interpreted as probabilities, then these two semiring interpretations can either find the “weight” of the parse with maximum weight or the total weight of all parse trees (a measure of the “total grammaticality” of a sentence). In this section, we give the specification of the weighted CKY algorithm [25], which is a dynamic programming algorithm for parsing using a context-free grammar in Chomsky normal form. The CKY algorithm is shown in Fig. 17. We show that fundamental algorithms for weighted (probabilistic) parsing can be derived as constrained **PRODUCT**s of CKY.

The unconstrained **PRODUCT** of CKY with itself (Fig. 18) is not inherently interesting. It is worth noting, however, that clause 43 there was a choice as to how to merge the c_1 and c_2 possibilities. The choice would not have existed if, instead of the presentation of CKY in Fig. 17, a common *binarized* variant of the algorithm, which introduces a new predicate in order to have at most two antecedents per Horn equation, had been fed to **PRODUCT**. The choice that we made in pairing was consistent with the choice that is forced in the binarized CKY program.

Product of Grammars. Fig. 19 describes a more interesting constrained version of Fig. 18. In particular, in all cases the constraints $I_1 = I_2, J_1 = J_2, K_1 = K_2, N_1 = N_2$ are added, so that instead of writing $c_{1\circ 2}(X_1, I_1, J_1, X_2, I_2, J_2)$ we just write $c_{1\circ 2}(X_1, X_2, I, J)$. This program simultaneously parses two different sentences using two different grammars, but

⁵ Chomsky normal form (CNF) means that the rules in the grammar are either binary with two nonterminals or unary with a terminal. We do not allow ϵ rules, which in general are allowed in CNF grammars.

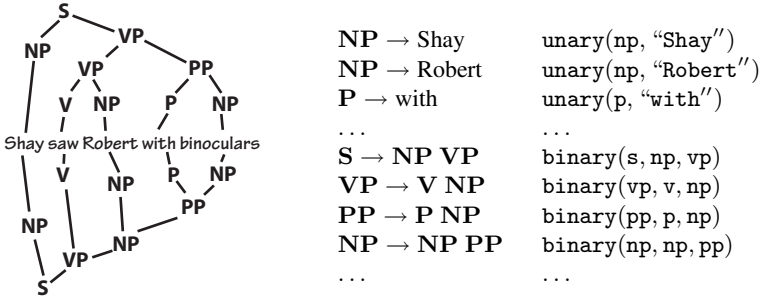


Fig. 16. An ambiguous sentence that can be parsed two ways in English (left), some of the Chomsky normal form rules for English grammar (center), and the corresponding axioms (right). There would also need to be five axioms of the form $\text{string}(\text{"Shay"}, 1)$, $\text{string}(\text{"saw"}, 2)$, etc.

$$\text{goal} \oplus= \text{start}(\text{S}) \otimes \text{length}(\text{N}) \otimes \text{c}(\text{S}, 0, \text{N}). \quad (36)$$

$$\text{c}(\text{X}, \text{I} - 1, \text{I}) \oplus= \text{unary}(\text{X}, \text{W}) \otimes \text{string}(\text{W}, \text{I}). \quad (37)$$

$$\text{c}(\text{X}, \text{I}, \text{K}) \oplus= \text{binary}(\text{X}, \text{Y}, \text{Z}) \otimes \text{c}(\text{Y}, \text{I}, \text{J}) \otimes \text{c}(\text{Z}, \text{J}, \text{K}). \quad (38)$$

Fig. 17. CKY: a weighted logic program implementing weighted CKY for algorithms involving weighted context free grammars in Chomsky normal form. Strictly speaking, CKY refers to a naïve bottom-up evaluation strategy for this program.

$$\text{goal}_{1 \otimes 2} \oplus= \text{length}_1(\text{N}_1) \otimes \text{length}_2(\text{N}_2) \otimes \quad (39)$$

$$\text{start}_1(\text{S}_1) \otimes \text{start}_2(\text{S}_2) \otimes \text{c}_{1 \otimes 2}(\text{S}_1, 0, \text{N}_1, \text{S}_2, 0, \text{N}_2).$$

$$\text{c}_{1 \otimes 2}(\text{X}_1, \text{I}_1 - 1, \text{I}_1, \text{X}_2, \text{I}_2 - 1, \text{I}_2) \oplus= \text{unary}_1(\text{X}_1, \text{W}_1) \otimes \text{string}_1(\text{W}_1, \text{I}_1) \otimes \quad (40)$$

$$\text{unary}_2(\text{X}_2, \text{W}_2) \otimes \text{string}_2(\text{W}_2, \text{I}_2).$$

$$\text{c}_{1 \otimes 2}(\text{X}_1, \text{I}_1 - 1, \text{I}_1, \text{X}_2, \text{I}_2, \text{K}_2) \oplus= \text{unary}_1(\text{X}_1, \text{W}_1) \otimes \text{string}_1(\text{W}_1, \text{I}_1) \otimes \quad (41)$$

$$\text{binary}_2(\text{X}_2, \text{Y}_2, \text{Z}_2) \otimes \text{c}_2(\text{Y}_2, \text{I}_2, \text{J}_2) \otimes \text{c}_2(\text{Z}_2, \text{J}_2, \text{K}_2).$$

$$\text{c}_{1 \otimes 2}(\text{X}_1, \text{I}_1, \text{K}_1, \text{X}_2, \text{I}_2 - 1, \text{I}_2) \oplus= \text{unary}_2(\text{X}_2, \text{W}_1) \otimes \text{string}_2(\text{W}_2, \text{I}_2) \otimes \quad (42)$$

$$\text{binary}_1(\text{X}_1, \text{Y}_1, \text{Z}_1) \otimes \text{c}_1(\text{Y}_1, \text{I}_1, \text{J}_1) \otimes \text{c}_1(\text{Z}_1, \text{J}_1, \text{K}_1).$$

$$\text{c}_{1 \otimes 2}(\text{X}_1, \text{I}_1, \text{K}_1, \text{X}_2, \text{I}_2, \text{K}_2) \oplus= \text{binary}_1(\text{X}_1, \text{Y}_1, \text{Z}_1) \otimes \text{binary}_2(\text{X}_2, \text{Y}_2, \text{Z}_2) \otimes \quad (43)$$

$$\text{c}_{1 \otimes 2}(\text{Y}_1, \text{I}_1, \text{J}_1, \text{Y}_2, \text{I}_2, \text{J}_2) \otimes \text{c}_{1 \otimes 2}(\text{Z}_1, \text{J}_1, \text{K}_1, \text{Z}_2, \text{K}_2, \text{J}_2).$$

Fig. 18. The full output of the PRODUCT transformation on two copies of CKY in Fig. 17

both parses must have the same *structure*. This constraint, then, can be compared to the constraints placed on the product of two graph-reachability programs to ensure that both paths have the same length.

Lexicalized CFG Parsing. An interesting variant of the previous rule involves *lexicalized grammars*, which are motivated in Fig. 20. Instead of describing a grammar using nonterminals denoting phrases (e.g., NP and VP), we can define a (context-free) *dependency grammar* [26] that encodes the syntax of a sentence in terms of parent-child relationships between words. In the case of the example of Fig. 20, the arrows below

$$\text{goal}_{1\circ 2} \oplus = \text{length}(N) \otimes \text{start}_1(S_1) \otimes \text{start}_2(S_2) \otimes c_{1\circ 2}(S_1, S_2, 0, N) \quad (44)$$

$$c_{1\circ 2}(X_1, X_2, I - 1, I) \oplus = \text{unary}_1(X_1, W_1) \otimes \text{string}_1(W_1, I) \otimes \text{unary}_2(X_2, W_2) \otimes \text{string}_2(W_2, I). \quad (45)$$

$$c_{1\circ 2}(X_1, X_2, I, K) \oplus = \text{binary}_1(X_1, Y_1, Z_1) \otimes \text{binary}_2(X_2, Y_2, Z_2) \otimes c_{1\circ 2}(Y_1, Y_2, I, J) \otimes c_{1\circ 2}(Z_1, Z_2, J, K). \quad (46)$$

Fig. 19. The program in Fig. 18 constrained to require parsing two different sentences with the same parse tree

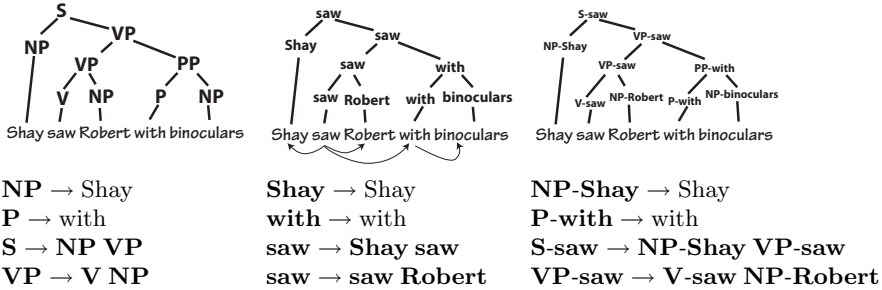


Fig. 20. On the left, the grammar previously shown. In the middle, a context-free *dependency grammar*, whose derivations can be seen as parse trees (above) or a set of dependencies (below). On the right, a lexicalized grammar. Sample rules are given for each grammar.

$$\text{goal}_{1\circ 2} \oplus = \text{length}(N) \otimes \text{start}_1(S_1) \otimes \text{start}_2(S_2) \otimes c_{1\circ 2}(S_1, S_2, 0, N) \quad (47)$$

$$c_{1\circ 2}(X_1, X_2, I - 1, I) \oplus = \text{unary}_1(X_1, W_1) \otimes \text{unary}_2(X_2, W_2) \otimes \text{string}(W, I). \quad (48)$$

$$c_{1\circ 2}(X_1, X_2, I, K) \oplus = \text{binary}_1(X_1, Y_1, Z_1) \otimes \text{binary}_2(X_2, Y_2, Z_2) \otimes c_{1\circ 2}(Y_1, Y_2, I, J) \otimes c_{1\circ 2}(Z_1, Z_2, J, K). \quad (49)$$

Fig. 21. Constraining Fig. 18 to simultaneously parse the *same* sentence with two grammars

the sentence in the middle establish “saw” as the root of the sentence; the word “saw” has three children (arguments and modifiers), one of which is the word “with,” which in turn has the child “binoculars.”

The simplest approach to describing a dependency grammar is to define it as a Chomsky normal form grammar where the nonterminal set is equivalent to the set of terminal symbols (so that the terminal “with” corresponds to a unique nonterminal **with**, and so on) and where all rules have the form $P \rightarrow P C$, $P \rightarrow C P$, and $W \rightarrow w$ (where X is the nonterminal version of terminal x).

Fig. 21 describes a further constrained program that, instead of parsing two unique strings in two different grammars with the same structure, parses a single string in two different grammars with the same structure. This new grammar recognizes a string if and only if both of the original grammars recognize it with isomorphic trees—a kind of “derivation intersection.” (This is not to be confused with intersections of context-free *languages*, which are not in general context-free languages [24].)

If we encode the regular grammar in the unary_1 and binary_1 relations and encode a dependency grammar in the unary_2 and binary_2 relations, then the product is a *lexicalized grammar*, like the third example from Fig. 20. In particular, it describes a lexicalized context-free grammar with a product of experts probability model [15], because the weight given to the production $\mathbf{A-X} \rightarrow \mathbf{B-X C-Y}$, for instance, is the semiring-product of the weight given to the production $\mathbf{A} \rightarrow \mathbf{B C}$ and the weight given to the dependency based production $\mathbf{X} \rightarrow \mathbf{X Y}$. If, instead of the axioms of the form $\text{binary}_1(X_1, Y_1, Z_1)$ and $\text{binary}_2(X_2, Y_2, Z_2)$ there were axioms of the form $\text{binary}_{1 \circ 2}(X_1, X_2, Y_1, Y_2, Z_1, Z_2)$ and clause 46 was changed accordingly, then the result would be a general lexicalized CKY [27].

Synchronous Parsing. Another extension to context-free parsing, *synchronous parsing*, can be derived using **PRODUCT** from two instances of CKY. Here two strings are parsed, each in a different alphabet with a different grammar (e.g., a French sentence and its English translation). A synchronous derivation consists of two trees and a correspondence between their nodes; different degrees of isomorphism may be imposed (e.g., in natural language, reordering is common, but dependencies tend to be mirrored through word-level translation). Constraining the **PRODUCT** of CKY with itself with side conditions to impose a one-to-one correspondence of nonterminals leads to a weighted logic program for a formalism known as inversion transduction grammar [28]. Constraining the **PRODUCT** of a more general CKY that includes empty unary rules “ $\mathbf{X} \rightarrow \epsilon$ ” and leaving rules that pair unary with binary antecedents removes the requirement that the sentences have the same length. In addition, the **PRODUCT** of the lexicalized CKY with itself leads to WLPs for more complex parsers like those described in [29] and [30] for more expressive formalisms.⁶

6 Conclusion

We have described a general framework for dynamic programming algorithms whose solutions correspond to proof values in two mutually constrained weighted logic programs. Our framework includes a program transformation, **PRODUCT**, which combines the two weighted logic programs that compute over two structures into a single weighted logic program for a joint proof. Appropriate constraints, encoded intuitively as variable unifications or side conditions in the WLP, are then added manually. The framework naturally captures many existing algorithms.

Acknowledgments

The authors acknowledge helpful comments from three anonymous ICLP reviewers, Jason Eisner, Frank Pfenning, David Smith, and Sylvia Rebolz. This research was supported by an NSF graduate fellowship to the second author and NSF grant IIS-0713265 and an IBM faculty award to the third author.

⁶ To derive the general synchronous parser described in [28] we have to perform another step of axiom generalization, described in [20]. In the current form, the axioms are factored as well.

References

1. Greco, S., Zaniolo, C.: Greedy algorithms in Datalog. *Theory Pract. Log. Program* 1(4), 381–407 (2001)
2. Ganzinger, H., McAllester, D.A.: Logical algorithms. In: Stuckey, P.J. (ed.) *ICLP 2002*. LNCS, vol. 2401, pp. 209–223. Springer, Heidelberg (2002)
3. Tarjan, R.E.: A unified approach to path problems. *Journal of the ACM* 28(3), 577–593 (1981)
4. Eisner, J., Goldlust, E., Smith, N.A.: Dyna: A declarative language for implementing dynamic programs. In: *Proc. of ACL (companion volume)* (2004)
5. Goodman, J.: Semiring parsing. *Computational Linguistics* 25(4), 573–605 (1999)
6. Shieber, S.M., Schabes, Y., Pereira, F.C.N.: Principles and implementation of deductive parsing. *Journal of Logic Programming* 24(1–2), 3–36 (1995)
7. Sikkel, K.: *Parsing Schemata: A Framework for Specification and Analysis of Parsing Algorithms*. Springer-Verlag New York, Inc., Secaucus, NJ, USA (1997)
8. McAllester, D.A.: On the complexity analysis of static analyses. In: Cortesi, A., Filé, G. (eds.) *SAS 1999*. LNCS, vol. 1694, pp. 312–329. Springer, Heidelberg (1999)
9. Klein, D., Manning, C.D.: Parsing and hypergraphs. *New developments in parsing technology*, 351–372 (2004)
10. Felzenszwalb, P.F., McAllester, D.: The generalized A* architecture. *Journal of Artificial Intelligence Research* 29, 153–190 (2007)
11. Eisner, J., Goldlust, E., Smith, N.A.: Compiling comp ling: practical weighted dynamic programming and the dyna language. In: *HLT 2005: Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing*, Morristown, NJ, USA, pp. 281–290. Association for Computational Linguistics (2005)
12. Huang, L., Chiang, D.: Better k-best parsing. In: *Proceedings of the Ninth International Workshop on Parsing Technologies (IWPT 2005)*, Vancouver, Canada (2005)
13. Eisner, J., Blatz, J.: Program transformations for optimization of parsing algorithms and other weighted logic programs. In: Wintner, S. (ed.) *Proceedings of FG 2006: The 11th Conference on Formal Grammar*, pp. 45–85. CSLI Publications (2007)
14. Hinton, G.E.: Training products of experts by minimizing contrastive divergence. *Neural Comput.* 14(8), 1771–1800 (2002)
15. Klein, D., Manning, C.D.: Fast exact inference with a factored model for natural language parsing. In: *Advances in Neural Information Processing Systems*, pp. 3–10. MIT Press, Cambridge (2002)
16. Liang, P., Klein, D., Jordan, M.: Agreement-based learning. In: Platt, J., Koller, D., Singer, Y., Roweis, S. (eds.) *Advances in Neural Information Processing Systems*, vol. 20, pp. 913–920. MIT Press, Cambridge (2008)
17. Chiang, D.: Hierarchical phrase-based translation. *Comput. Linguist.* 33(2), 201–228 (2007)
18. Cohen, S.B., Smith, N.A.: Joint morphological and syntactic disambiguation. In: *Proceedings of EMNLP-CoNLL 2007*, pp. 208–217 (2007)
19. Sutton, C., McCallum, A.: Piecewise training for undirected models. In: *Proceedings of the 21th Annual Conference on Uncertainty in Artificial Intelligence (UAI 2005)*, Arlington, Virginia, p. 568. AUAI Press (2005)
20. Cohen, S.B., Simmons, R.J., Smith, N.A.: Products of weighted logic programs. Technical Report CMU-LTI-08-009, Carnegie Mellon University (2008)
21. Levenshtein, V.: Binary codes capable of correcting spurious insertions and deletions of ones. *Problems of Information Transmission* 1, 8–17 (1965)
22. Mohri, M.: Finite-state transducers in language and speech processing. *Comput. Linguist.* 23(2), 269–311 (1997)

23. Manning, C., Schütze, H.: Foundations of Statistical Natural Language Processing. MIT Press, Cambridge (1999)
24. Hopcroft, J.E., Ullman, J.D.: Introduction to Automata Theory, Languages, and Computation. Addison-Wesley, Reading (1979)
25. Cocke, J., Schwartz, J.T.: Programming languages and their compilers: Preliminary notes. Technical report, Courant Institute of Mathematical Sciences, New York University (1970)
26. Gaifman, H.: Dependency systems and phrase-structure systems. *Information and Control* 8 (1965)
27. Eisner, J., Satta, G.: Efficient parsing for bilexical context-free grammars and head automaton grammars. In: Proceedings of the 37th annual meeting of the Association for Computational Linguistics on Computational Linguistics, Morristown, NJ, USA, pp. 457–464. Association for Computational Linguistics (1999)
28. Wu, D.: Stochastic inversion transduction grammars and bilingual parsing of parallel corpora. *Computational Linguistics* 23(3), 377–404 (1997)
29. Melamed, I.D.: Multitext grammars and synchronous parsers. In: NAACL 2003: Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology, Morristown, NJ, USA, pp. 79–86. Association for Computational Linguistics (2003)
30. Zhang, H., Gildea, D.: Stochastic lexicalized inversion transduction grammar for alignment. In: ACL 2005: Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics, Morristown, NJ, USA, pp. 475–482. Association for Computational Linguistics (2005)