Proceedings of the Twenty-Fourth International Florida Artificial Intelligence Research Society Conference

A Default Logical Semantics for Defeasible Argumentation

Gabriele Kern-Isberner
Department of Computer Science
University of Technology Dortmund, Germany

Guillermo R. Simari
Department of Computer Science & Engineering
Universidad Nacional del Sur, Argentina

Abstract

Defeasible argumentation and default reasoning are usually perceived as two similar, but distinct approaches to commonsense reasoning. In this paper, we combine these two fields by viewing (defeasible resp. default) rules as a common crucial part in both areas. We will make use of possible worlds semantics from default reasoning to provide examples for arguments, and carry over the notion of plausibility to the argumentative framework. Moreover, we base a priority relation between arguments on the tolerance partitioning of system Z and obtain a criterion phrased in system Z terms that ensures warrancy in defeasible argumentation.

Key Words: Defeasible argumentation, default reasoning, system Z, possible worlds

1 Introduction

Argumentation techniques have gained increasing relevance in artificial intelligence research during the past two decades, mainly due to their wide range of applicability within agent systems, providing bases both for single agent reasoning and deliberation, and for coordination and exchange of information in multiagent systems. Basically, argumentation investigates how arguments attack and defend one another, thereby pondering in a dialectical way the pro and cons of options to make up a proper decision. Presumably the most foundational and abstract framework of argumentation has been provided by Dung (Dung 1995), but lots of other more constructive approaches to argumentation have been brought forward as well, making use of logical elements like rules and deduction, and of priority relations to decide between good and worse arguments (see, e.g., (Besnard and Hunter 2008)).

*Defeasible Logic Programming (*DeLP*)* (García and Simari 2004) combines logic programming with defeasible argumentation, allowing the representation of tentative knowledge and leaving for the argumentation mechanism the task of finding those conclusions for which it is not possible to find valid reasons against. DeLP works in a highly dialectical way, allowing series of attacks and counterattacks to finally mark those statements as *warranted* for which all attackers could be invalidated. Although DeLP rests on logic programming grounds and its warrant procedure corresponds nicely to human argumentation behaviour, it has not yet been possible to characterize its semantics of warrant along established formal logical lines (Thimm and Kern-Isberner 2008).

In this paper, we take an alternative semantical way to support the conclusions drawn by DeLP. First, we introduce possible worlds as examples and counterexamples of arguments, and characterize crucial notions like *attack* and *concordance* in DeLP in these terms. Then, we make use of ordinal conditional functions (Spohn 1988) that provide degrees of (dis)belief to possible worlds in order to assign priority information to arguments. The basic idea here is that arguments are as convincing and successful as their most plausible examples. So, arguments with more plausible examples should prevail, in particular when compared to counterarguments and counterexamples.

Then, in order to sharpen our results, we make use of the distinguished system Z approach (Goldszmidt and Pearl 1996) as a particularly well-behaved ordinal conditional function to assign plausibility to examples of arguments and prioritize arguments by Z-specificity values. We prove a sufficient condition to ensure warrant in DeLP, and show that our combined Z-DeLP-approach is able to solve the so-called *drowning problem*.

The rest of this paper is organized as follows: Section 2 recalls the background needed for ordinal conditional functions, system Z and DeLP. Then, our semantics based on examples of arguments is developed in section 3, and is connected to plausibility in section 4. Afterwards, section 5 presents the combined Z-DeLP-approach, and section 6 concludes this paper. All notions and techniques are illustrated by a running example.

2 Formal and methodological background

Logical prerequisites. Let $\mathcal{L}$ be a finitely generated propositional language with atoms $a, b, c, \ldots$, and with formulas $A, B, C, \ldots$. For conciseness of notation, we will omit the logical *and*-connector, writing AB instead of $A \wedge B$, and overlining formulas will indicate negation, i.e. $\overline{A}$ means $\neg A$. Let Ω denote the set of possible worlds over $\mathcal{L}$; Ω will be taken here simply as the set of all propositional interpretations over $\mathcal{L}$. $\omega \models A$ means that the propositional formula

$A \in \mathcal{L}$ holds in the possible world $\omega \in \Omega$. The set of models of A is denoted by $\mathit{Mod}(A) = \{\omega \in \Omega \mid \omega \models A\}$.

By making use of a new binary operator $|$, we obtain the set $(\mathcal{L} \mid \mathcal{L}) = \{(B|A) \mid A, B \in \mathcal{L}\}$ of conditionals over $\mathcal{L}$. In the framework of default reasoning, $(B|A)$ formalizes "*if A then B is plausible*" and establishes a defeasible connection between the *antecedent* A and the *consequent* B. In the framework of argumentation, $(B|A)$ may be read as "*If there are reasons to believe A, then there are reasons to believe B*". Here, conditionals are supposed not to be nested, that is, antecedent and consequent of a conditional will be propositional formulas. A conditional $(B|A)$ is an object of a three-valued nature, partitioning the set of worlds Ω in three parts: those worlds satisfying AB and thus *verifying* the conditional, those worlds satisfying $A\overline{B}$, thus *falsifying* the conditional, and those worlds not fulfilling the premise A and so which the conditional may not be applied to at all (cf. (DeFinetti 1974)).

OCFs and system Z. *Ordinal conditional functions* (*OCFs*, also called *ranking functions*) are functions $\kappa : \Omega \to \mathbb{N} \cup \{\infty\}$ with $\kappa^{-1}(0) \neq \emptyset$, expressing degrees of plausibility of propositional formulas A by specifying degrees of disbeliefs of their negations $\overline{A}$ (cf. (Spohn 1988; Goldszmidt and Pearl 1996)). More formally, we have $\kappa(A) := \min\{\kappa(\omega) \mid \omega \models A\}$, so that $\kappa(A \vee B) = \min\{\kappa(A), \kappa(B)\}$. A conditional $(B|A)$ is accepted in the epistemic state represented by κ, or κ *satisfies* $(B|A)$, written as $\kappa \models (B|A)$, iff $\kappa(AB) < \kappa(A\overline{B})$, i.e. iff AB is more plausible than $A\overline{B}$. A set $\Delta = \{\delta_i = (B_i|A_i) \mid 1 \leq i \leq n\}$ of conditionals is said to be *consistent* iff there exists an OCF κ that accepts all conditionals in Δ, i.e. for which holds $\kappa(A_iB_i) < \kappa(A_i\overline{B_i}), 1 \leq i \leq n$. Consistency can be checked easily by applying the notion of *tolerance*. A conditional $(B|A)$ is *tolerated* by a set of conditionals Δ iff there is a world ω such that ω verifies $(B|A)$ and ω does not falsify any of the conditionals in Δ. Δ is consistent, iff there is an ordered partition $\Delta_0, \Delta_1, \ldots, \Delta_k$ of Δ such that each conditional in Δ_m is tolerated by $\bigcup_{j=m}^{k} \Delta_j$, $0 \leq m \leq k$ (Goldszmidt and Pearl 1996).

A well-known method to compute an ordinal conditional function which accepts all conditionals in a (finite) set $\Delta = \{\delta_i = (B_i|A_i) \mid 1 \leq i \leq n\}$ is *system* Z of Goldszmidt and Pearl ((Goldszmidt and Pearl 1996)). System Z makes use of the tolerance partitioning described above, where each subset Δ_i of the partition is maximal, i.e. contains as many conditionals as possible. Then to each $\delta \in \Delta$, a rank is associated by setting $Z(\delta) = j$ iff $\delta \in \Delta_j$. The system Z ranking function, κ_z, accepting all conditionals in Δ is then given by

$$\kappa_z(\omega) = \begin{cases} 0, \text{ if } \omega \text{ does not falsify any } \delta_i, \\ 1 + \max\limits_{\substack{1 \leq i \leq n \\ \omega \models A_i\overline{B}_i}} Z(\delta_i), \text{otherwise} \end{cases} \tag{1}$$

κ_z assigns to each world ω the lowest possible rank admissible with respect to the constraints in Δ. System Z can be used to define a nonmonotonic inference relation $\mid\!\sim_Z$: Given some consistent set of conditionals Δ and propositional formulas A, B, we have $A \mid\!\sim_Z B$ iff $\kappa_z \models (B|A)$, i.e., iff in the context of A, B turns out to be plausible. For more details see, for instance, (Goldszmidt and Pearl 1996). We illustrate default reasoning with system Z in the following benchmark example taken from (García and Simari 2004).

Example 1 We consider the propositional variables b *bird*, p *penguin*, c *chicken*, s *is_scared*, f *flies*, w *has_wings*, and the set of conditionals: $\Delta = \{\delta_1 = (b|c), \delta_2 = (b|p), \delta_3 = (f|b), \delta_4 = (\overline{f}|p), \delta_5 = (\overline{f}|c), \delta_6 = (f|cs), \delta_7 = (w|b)\}$. The tolerance partitioning used by system Z is $\Delta_0 = \{\delta_3, \delta_7\}, \Delta_1 = \{\delta_1, \delta_2, \delta_4, \delta_5\}, \Delta_2 = \{\delta_6\}$. We compute $\kappa_z(pbf) = 2 > 1 = \kappa_z(pb\overline{f})$, so penguin-birds do not fly, as expected. Also scared penguins do not fly, since $\kappa_z(psf) = 2 > 1 = \kappa_z(ps\overline{f})$.

The basic problem of default reasoning is to decide if some formula A is a plausible consequence of a set of conditionals, Δ, given some prerequisite B. This amounts to checking if the conditional $(B|A)$ is accepted on the ground of Δ. Ranking functions and, in particular, system Z provide a convenient semantics to verify this inference relation.

Defeasible argumentation. We will present the basics of the formalism of DeLP, following the presentation in (García and Simari 2004). A *defeasible logic program (de.l.p.)* $\mathcal{P} = (\Pi, \Delta)$ consists of a set Π of facts and strict logic rules, and a set Δ of defeasible rules which are written as conditionals $(L|B_1 \ldots B_n)$ with literals $L, B_1, \ldots, B_n$ in this paper. Let $\Delta' \subseteq \Delta$ be a subset of Δ. A literal L can be *defeasibly derived* from Δ', iff there exists a finite sequence $L_1, \ldots, L_n = L$ of ground literals, such that each L_i is either a fact in Π or there exists a rule in $\Pi \cup \Delta'$ with head L_i and body $\{B_1, \ldots, B_m\}$, and every literal B_j in the body is such that $B_j \in \{L_k\}_{k<i}$. $\Pi \cup \Delta'$ is called *contradictory* iff there is a literal L such that both L and $\overline{L}$ have defeasible derivations from $\Pi \cup \Delta'$. For any *de.l.p.* $\mathcal{P}$ we will presuppose that Π is non-contradictory.

Given a *de.l.p.* $\mathcal{P} = (\Pi, \Delta)$ and a literal L, $\mathcal{A}$ is an argument for L, denoted $\langle \mathcal{A}, L \rangle$, if $\mathcal{A}$ is a set of defeasible rules in Δ such that:

1. there exists a defeasible derivation of L from $\Pi \cup \mathcal{A}$;
2. $\Pi \cup \mathcal{A}$ is non-contradictory;
3. there is no $\mathcal{A}' \subsetneq \mathcal{A}$ such that $\mathcal{A}'$ satisfies (1) and (2), *i.e.*, in that sense $\mathcal{A}$ is minimal.

An argument $\langle \mathcal{B}, Q \rangle$ is a sub-argument of $\langle \mathcal{A}, L \rangle$ if $\mathcal{B}$ is subset of $\mathcal{A}$.

Argument $\langle \mathcal{A}_1, L_1 \rangle$ *attacks*, or *counterargues* another $\langle \mathcal{A}_2, L_2 \rangle$ at a literal L if there exists a sub-argument of $\langle \mathcal{A}_2, L_2 \rangle$, $\langle \mathcal{A}, L \rangle$, *i.e.*, $\mathcal{A} \subseteq \mathcal{A}_2$, such that there exists a literal P verifying both $\Pi \cup \{L, L_1\} \vdash P$ and $\Pi \cup \{L, L_1\} \vdash \overline{P}$. Note that an argument $\langle \emptyset, L \rangle$ with $L \in \Pi$ can not be attacked since all arguments have to be consistent with Π.

Example 2 We will look at example 1 in the DeLP-framework. Let propositionals and conditionals be given as specified in example 1. Consider the defeasible logic program $\mathcal{P}_1 = (\{cs\}, \Delta)$. Then the following arguments can be built supporting f resp. $\overline{f}$:

$$\langle \mathcal{A}_1, f\rangle, \qquad \mathcal{A}_1 = \{(b|c), (f|b)\};$$
$$\langle \mathcal{A}_2, \overline{f}\rangle, \qquad \mathcal{A}_2 = \{(\overline{f}|c)\};$$
$$\langle \mathcal{A}_3, f\rangle, \qquad \mathcal{A}_3 = \{(f|cs)\}.$$

Clearly, $\langle \mathcal{A}_2, \overline{f}\rangle$ attacks $\langle \mathcal{A}_1, f\rangle$, and $\langle \mathcal{A}_3, f\rangle$ attacks $\langle \mathcal{A}_2, \overline{f}\rangle$.

An argumentation process proceeds through comparisons among arguments. The standard criterion of comparison used in DELP is *specificity*, but any criteria that establishes preference could be used. Since we will use other preference criteria here, we will not go into further details here.

If $\langle \mathcal{A}_1, L_1\rangle$ and $\langle \mathcal{A}_2, L_2\rangle$ are two arguments $\langle \mathcal{A}_1, L_1\rangle$ is a *proper defeater* for $\langle \mathcal{A}_2, L_2\rangle$ at literal L iff there exists a sub-argument of $\langle \mathcal{A}_2, L_2\rangle$, $\langle \mathcal{A}, L\rangle$ such that $\langle \mathcal{A}_1, L_1\rangle$ counterargues $\langle \mathcal{A}_2, L_2\rangle$ at L and $\langle \mathcal{A}_1, L_1\rangle$ is strictly preferred over $\langle \mathcal{A}, L\rangle$. Alternatively, $\langle \mathcal{A}_1, L_1\rangle$ is a *blocking defeater* for $\langle \mathcal{A}_2, L_2\rangle$ at literal L iff there exists a sub-argument of $\langle \mathcal{A}_2, L_2\rangle$, $\langle \mathcal{A}, L\rangle$ such that $\langle \mathcal{A}_1, L_1\rangle$ counterargues $\langle \mathcal{A}_2, L_2\rangle$ at L and neither $\langle \mathcal{A}_1, L_1\rangle$ is strictly preferred over $\langle \mathcal{A}, L\rangle$ nor is $\langle \mathcal{A}, L\rangle$ preferred over $\langle \mathcal{A}_1, L_1\rangle$. If $\langle \mathcal{A}_1, L_1\rangle$ is either a proper or a blocking defeater of $\langle \mathcal{A}_2, L_2\rangle$, it is said to be a *defeater* of the latter.

An *argumentation line* for an argument $\langle \mathcal{A}_0, L_0\rangle$ is a sequence $\lambda = [\langle \mathcal{A}_0, L_0\rangle, \langle \mathcal{A}_1, L_1\rangle, \langle \mathcal{A}_2, L_2\rangle, \cdots]$ where for each $i > 0$ $\langle \mathcal{A}_{i+1}, L_{i+1}\rangle$ is a defeater of $\langle \mathcal{A}_i, L_i\rangle$. $\lambda_S = [\langle \mathcal{A}_0, L_0\rangle, \langle \mathcal{A}_2, L_2\rangle, \langle \mathcal{A}_4, L_4\rangle, \cdots]$ is the sequence of *supporting* argument of λ, while the sequence of *interfering* ones is $\lambda_I = [\langle \mathcal{A}_1, L_1\rangle, \langle \mathcal{A}_3, L_3\rangle, \langle \mathcal{A}_5, L_5\rangle, \cdots]$.

An *acceptable* argumentation line in a defeasible program $\mathcal{P} = (\Pi, \Delta)$ is a finite sequence $\lambda = [\langle \mathcal{A}_0, L_0\rangle, \cdots, \langle \mathcal{A}_n, L_n\rangle]$ such that some constraints on the addition of arguments to the sequence are considered. For example, not allowing the addition of an already introduced argument, or using an argument simultaneously as a supporting and as an interfering argument (see (García and Simari 2004) for an in-depth analysis of the problem). A constraint that we will postulate for any acceptable argumentation line to hold is *concordance* of all supporting respectively interfering arguments.

Definition 3 (concordance) *A set of arguments $\mathcal{A}_i, 1 \leq i \leq m$, of a defeasible logic program (Π, Δ) is called* concordant *iff $\Pi \cup \bigcup_{i=1}^{n} \mathcal{A}_i$ is non-contradictory.*

To answer a query Q, the *warrant procedure* builds up a candidate argument $\langle \mathcal{A}, Q\rangle$. Then, it associates to this argument a *dialectical tree* $\mathcal{T}_{\langle \mathcal{A}, Q\rangle}$ as follows:

1. The root of the tree is labeled, $\langle \mathcal{A}_0, Q_0\rangle$, *i.e.*, $\mathcal{A}_0 = \mathcal{A}$ and $Q_0 = Q$.
2. Let n be a non-root node, with label $\langle \mathcal{A}_n, Q_n\rangle$ and $\lambda = [\langle \mathcal{A}_0, Q_0\rangle, \cdots, \langle \mathcal{A}_n, Q_n\rangle]$ the labels in the path from the root to n. Let $\mathbf{B} = \{\langle \mathcal{B}_1, R_1\rangle, \cdots, \langle \mathcal{B}_k, R_k\rangle\}$ be the set of all the defeaters for $\langle \mathcal{A}_n, Q_n\rangle$. For $1 \leq i \leq k$, if $\lambda' = [\langle \mathcal{A}_0, Q_0\rangle, \cdots, \langle \mathcal{A}_n, Q_n\rangle, \langle \mathcal{B}_i, R_i\rangle]$ is an acceptable argumentation line, n has a child n_i labeled $\langle \mathcal{B}_i, R_i\rangle$. If $\mathbf{B} = \emptyset$ or no $\langle \mathcal{B}_i, R_i\rangle \in \mathcal{B}$ is such that λ' is acceptable, then n is a leaf of the tree.

The nodes of $\mathcal{T}_{\langle \mathcal{A}, Q\rangle}$ can be marked U (undefeated) or D (defeated), yielding a tagged tree $\mathcal{T}^*_{\langle \mathcal{A}, Q\rangle}$ as follows:

- All leaves of $\mathcal{T}_{\langle \mathcal{A}, Q\rangle}$ are marked U in $\mathcal{T}^*_{\langle \mathcal{A}, Q\rangle}$.
- If $\langle \mathcal{B}, R\rangle$ is the label of a node which is not a leaf, the node will be marked U in $\mathcal{T}^*_{\langle \mathcal{A}, Q\rangle}$ if every child is marked D. Otherwise, if at least one of its children is marked U, it is marked as D.

Then, given an argument $\langle \mathcal{A}, Q\rangle$ and its associated tagged tree $\mathcal{T}^*_{\langle \mathcal{A}, Q\rangle}$, if the root is marked U, the literal Q is said to be *warranted*, and $\mathcal{A}$ is said to be the *warrant* for Q.

In this paper, we will be assuming that there are no strict rules, so that the strict logical knowledge is represented by facts, and will write Φ instead of Π. With this prerequisite, argument $\langle \mathcal{A}_1, L_1\rangle$ attacks another $\langle \mathcal{A}_2, L_2\rangle$ at a literal L if there exists a sub-argument $\langle \mathcal{A}, L\rangle$ of $\langle \mathcal{A}_2, L_2\rangle$ such that $L = \overline{L_1}$. Concordance can be checked easily, too.

Proposition 4 *Let $\mathcal{P} = (\Pi, \Delta)$ be a defeasible logic program the strict part Π of which contains only facts. A set of arguments $\mathcal{A}_1, \ldots, \mathcal{A}_m$ of $\mathcal{P}$ is concordant iff for any two rules $\delta_1 \in \mathcal{A}_i, \delta_2 \in \mathcal{A}_j$, it holds that $head(\delta_1) \neq \overline{head(\delta_2)}$.*

Defeasible rules can be taken as conditionals, and we will use both terms synonymously. In the rest of the paper, a defeasible logic programm $\mathcal{P} = (\Phi, \Delta)$ consists of a collection of facts, Φ, and conditionals, Δ, both specified in a suitable propositional language $\mathcal{L}$ with Ω being the appertaining set of possible worlds. As a further important prerequisite, we postulate that the set Δ of conditionals be consistent (as defined in section 2). Then, this setting is fully compatible to the methodology of OCF's and system Z described in section 2, and we will elaborate this connection in the following sections. For ease of notation, Φ will also denote the conjunction of all facts of $\mathcal{P}$.

3 A semantics of examples and counterexamples for arguments

The basic idea of argumentation is to support beliefs by arguments that are often built from rules. This logical underpinning of arguments is crucial for DELP, but can not be taken into regard by Dung's abstract argumentation framework (Dung 1995) which is considered to be one of the standard approaches to give semantics to arguments. Here, we propose a semantical approach to arguments which works particularly well if the arguments have a logical structure but might be applied more generally. Our approach differs from Dung's framework and also from more logical approaches like those presented in (Besnard and Hunter 2008). The basic idea first to be developed in this paper is that arguments are as convincing as their (most plausible) examples. We will introduce the notion of examples and counterexamples of arguments in this section to justify a priority relation between arguments. Plausibility issues will be taken into considereation in the next section.

Possible worlds verifying all conditionals occurring in an argument may be used as examples to illustrate the statement of this argument.

Definition 5 (Examples, counterexamples) *Let $\mathcal{P} = (\Phi, \Delta)$ be a defeasible logic program. Let $\omega \in \Omega$ be a possible world, and let $\langle \mathcal{A}, L\rangle$ be an argument in $\mathcal{P}$.*

ω is an example for $\langle \mathcal{A}, L\rangle$ *iff ω satisfies all facts, $\omega \models \Phi$, and ω verifies all rules in $\mathcal{A}$. ω is a* counterexample to $\langle \mathcal{A}, L\rangle$ *iff $\omega \models \Phi$ and there is at least one rule in $\mathcal{A}$ that is falsified by ω. ω is a* supported counterexample to $\langle \mathcal{A}, L\rangle$ *iff ω is a counterexample to $\langle \mathcal{A}, L\rangle$ and there is an argument $\langle \mathcal{A}', L'\rangle$ such that ω is an example of $\langle \mathcal{A}', L'\rangle$.*

The set of examples of an argument $\langle \mathcal{A}, L\rangle$ is denoted by $\langle \mathcal{A}, L\rangle^+$, the set of counterexamples by $\langle \mathcal{A}, L\rangle^-$.

Proposition 6 *Every argument $\langle \mathcal{A}, L\rangle$ has examples; more precisely*[1],

$$\langle \mathcal{A}, L\rangle^+ = \mathit{Mod}(\Phi \wedge \bigwedge_{\delta \in \mathcal{A}} \mathit{head}(\delta)).$$

The set of counterexamples to $\langle \mathcal{A}, L\rangle$ is given by

$$\langle \mathcal{A}, L\rangle^- = \mathit{Mod}(\Phi \wedge \bigvee_{\delta \in \mathcal{A}} \overline{\mathit{head}(\delta)}).$$

Proof: Since $\Phi \cup \mathcal{A}$ is not contradictory, $\Phi \wedge \bigwedge_{\delta \in \mathcal{A}} \mathit{head}(\delta)$ is satisfiable, hence $\mathit{Mod}(\Phi \wedge \bigwedge_{\delta \in \mathcal{A}} \mathit{head}(\delta)) \neq \emptyset$. Let $\omega \in \mathit{Mod}(\Phi \wedge \bigwedge_{\delta \in \mathcal{A}} \mathit{head}(\delta))$. Then $\omega \models \Phi$, and, since all conditionals in $\mathcal{A}$ are applicable when the preceding conditionals have been applied, and ω satisfies all heads of conditionals in $\mathcal{A}$, ω verifies all conditionals occurring in $\mathcal{A}$. So, $\omega \in \langle \mathcal{A}, L\rangle^+$.

Conversely, let $\omega \in \langle \mathcal{A}, L\rangle^+$; then $\omega \models \Phi$ and ω verifies all conditionals in $\mathcal{A}$, in particular, ω satisfies all heads of rules occurring in $\mathcal{A}$. Hence $\omega \in \mathit{Mod}(\Phi \wedge \bigwedge_{\delta \in \mathcal{A}} \mathit{head}(\delta))$.

For the counterexamples, it is obvious that $\langle \mathcal{A}, L\rangle^- \subseteq \mathit{Mod}(\Phi \wedge \bigvee_{\delta \in \mathcal{A}} \overline{\mathit{head}(\delta)})$; in particular, if $\mathcal{A} = \emptyset$ then $\langle \mathcal{A}, L\rangle^- = \emptyset = \mathit{Mod}(\Phi \wedge \bigvee_{\delta \in \mathcal{A}} \overline{\mathit{head}(\delta)})$. Conversely, if $\omega \in \mathit{Mod}(\Phi \wedge \bigvee_{\delta \in \mathcal{A}} \overline{\mathit{head}(\delta)})$, then ω must falsify a (first) rule in $\mathcal{A}$, so $\omega \in \langle \mathcal{A}, L\rangle^-$. □

Hence, the models of the facts Φ provide a common reservoir of examples and counterexamples for all arguments of the *de.l.p.* $\mathcal{P}$:

Corollary 7 *Let $\langle \mathcal{A}, L\rangle$ be an argument. Then $\langle \mathcal{A}, L\rangle^+ \cup \langle \mathcal{A}, L\rangle^- = \mathit{Mod}(\Phi)$.*

Example 8 For the arguments $\langle \mathcal{A}_1, f\rangle, \langle \mathcal{A}_2, \overline{f}\rangle, \langle \mathcal{A}_3, f\rangle$ stated in example 2, examples and counterexamples are given as follows:

$$\begin{array}{ll} \langle \mathcal{A}_1, f\rangle^+ = \mathit{Mod}(csbf) & \langle \mathcal{A}_1, f\rangle^- = \mathit{Mod}(cs(\overline{b} \vee \overline{f})) \\ \langle \mathcal{A}_2, \overline{f}\rangle^+ = \mathit{Mod}(cs\overline{f}) & \langle \mathcal{A}_2, \overline{f}\rangle^- = \mathit{Mod}(csf) \\ \langle \mathcal{A}_3, f\rangle^+ = \mathit{Mod}(csf) & \langle \mathcal{A}_3, f\rangle^- = \mathit{Mod}(cs\overline{f}) \end{array}$$

Hence, $\omega_1 = csb\overline{p}fw$ is an example of $\langle \mathcal{A}_1, f\rangle$ and $\langle \mathcal{A}_3, f\rangle$ and a counterexample of $\langle \mathcal{A}_2, \overline{f}\rangle$. Reciprocally, $\omega_2 = csb\overline{p}\overline{f}w$ is an example of $\langle \mathcal{A}_2, \overline{f}\rangle$, and a counterexample of $\langle \mathcal{A}_1, f\rangle$ and $\langle \mathcal{A}_3, f\rangle$.

Relationships between arguments can be characterized easily by considering their examples:

[1] As usual, we set $\bigwedge \emptyset = \top$ and $\bigvee \emptyset = \bot$.

Proposition 9 *Let $\langle \mathcal{A}_1, L_1\rangle, \langle \mathcal{A}_2, L_2\rangle$ be two arguments.*

If $\langle \mathcal{A}_1, L_1\rangle$ attacks $\langle \mathcal{A}_2, L_2\rangle$, then all examples of $\langle \mathcal{A}_1, L_1\rangle$ are (supported) counterexamples to $\langle \mathcal{A}_2, L_2\rangle$, i.e. $\langle \mathcal{A}_1, L_1\rangle^+ \subseteq \langle \mathcal{A}_2, L_2\rangle^-$.

Conversely, if all examples of $\langle \mathcal{A}_1, L_1\rangle$ are counterexamples to $\langle \mathcal{A}_2, L_2\rangle$, then there is a sub-argument of $\langle \mathcal{A}_1, L_1\rangle$ that attacks $\langle \mathcal{A}_2, L_2\rangle$.

Proof: First, assume that $\langle \mathcal{A}_1, L_1\rangle$ attacks $\langle \mathcal{A}_2, L_2\rangle$ at literal L. Then there exists a sub-argument $\langle \mathcal{A}, L\rangle$ of $\langle \mathcal{A}_2, L_2\rangle$ such that $L = \overline{L_1}$. Any example ω of $\langle \mathcal{A}_1, L_1\rangle$ must verify any rule occurring in $\mathcal{A}_1$, in particular, $\omega \models L_1 = \overline{L}$, and L is the head of a rule occurring in $\mathcal{A}_2$. By Prop. 6, ω is a counterexample to $\langle \mathcal{A}_2, L_2\rangle$. Hence, $\langle \mathcal{A}_1, L_1\rangle^+ \subseteq \langle \mathcal{A}_2, L_2\rangle^-$.

Assume now that $\langle \mathcal{A}_1, L_1\rangle^+ \subseteq \langle \mathcal{A}_2, L_2\rangle^-$ holds. Since the set of examples and counterexamples of an argument are disjoint, we have $\langle \mathcal{A}_1, L_1\rangle^+ \cap \langle \mathcal{A}_2, L_2\rangle^+ = \emptyset$. By Prop. 6, $\Phi \wedge \bigwedge_{\delta \in \mathcal{A}_1} \mathit{head}(\delta) \wedge \bigwedge_{\delta \in \mathcal{A}_2} \mathit{head}(\delta)$ must be inconsistent. So, there are rules $\delta_1 \in \mathcal{A}_1, \delta_2 \in \mathcal{A}_2$ with $\mathit{head}(\delta_1) = Q_1, \mathit{head}(\delta_2) = Q_2$ such that $Q_1 = \overline{Q_2}$. Let $\langle \mathcal{A}_1', Q_1\rangle$ the subargument of $\langle \mathcal{A}_1, L_1\rangle$ that ends in Q_1. Then $\langle \mathcal{A}_1', Q_1\rangle$ attacks $\langle \mathcal{A}_2, L_2\rangle$ at Q_2. □

Proposition 10 *A set of arguments $\langle \mathcal{A}_i, L_i\rangle, 1 \leq i \leq m$, is concordant iff they have common examples, i.e. iff $\bigcap_{1 \leq i \leq m} \langle \mathcal{A}_i, L_i\rangle^+ \neq \emptyset$.*

Proof: By Prop. 4, a set of arguments $\mathcal{A}_i$, $1 \leq i \leq m$, is concordant iff there are no rules with conflicting heads in $\bigcup_{1 \leq i \leq m} \mathcal{A}_i$. Hence, concordance holds iff $\mathit{Mod}(\Phi \wedge \bigwedge_{\delta \in \bigcup_{1 \leq i \leq m} \mathcal{A}_i} \mathit{head}(\delta)) \neq \emptyset$. Since $\bigcap_{1 \leq i \leq m} \langle \mathcal{A}_i, L_i\rangle^+ = \bigcap_{1 \leq i \leq m} \mathit{Mod}(\Phi \wedge \bigwedge_{\delta \in \mathcal{A}_i} \mathit{head}(\delta)) = \mathit{Mod}(\Phi \wedge \bigwedge_{\delta \in \bigcup_{1 \leq i \leq m} \mathcal{A}_i} \mathit{head}(\delta))$, the proposition is proved. □

Corollary 11 *Two arguments $\langle \mathcal{A}_1, L_1\rangle$ and $\langle \mathcal{A}_2, L_2\rangle$ are either concordant, or there are subarguments of $\langle \mathcal{A}_1, L_1\rangle$ and $\langle \mathcal{A}_2, L_2\rangle$ that attack the other argument.*

Proof: For any two arguments $\langle \mathcal{A}_1, L_1\rangle$ and $\langle \mathcal{A}_2, L_2\rangle$ either $\langle \mathcal{A}_1, L_1\rangle^+ \cap \langle \mathcal{A}_2, L_2\rangle^+ = \emptyset$, or $\langle \mathcal{A}_1, L_1\rangle^+ \cap \langle \mathcal{A}_2, L_2\rangle^+ \neq \emptyset$. In the latter case, they are concordant, by Proposition 10. In the former case, we have $\mathcal{A}_1{}^+ \subseteq \mathcal{A}_2{}^-$ and $\mathcal{A}_2{}^+ \subseteq \mathcal{A}_1{}^-$. By Proposition 9, there are subarguments of $\langle \mathcal{A}_1, L_1\rangle$ and $\langle \mathcal{A}_2, L_2\rangle$ that attack the other argument. □

Moreover, the subargument relationship has also implications for the examples.

Proposition 12 *If $\langle \mathcal{A}_1, L_1\rangle$ is a subargument of $\langle \mathcal{A}_2, L_2\rangle$, then all examples of $\langle \mathcal{A}_2, L_2\rangle$ are also examples of $\langle \mathcal{A}_1, L_1\rangle$, i.e. $\langle \mathcal{A}_2, L_2\rangle^+ \subseteq \langle \mathcal{A}_1, L_1\rangle^+$.*

Proof: If $\langle \mathcal{A}_1, L_1\rangle$ is a subargument of $\langle \mathcal{A}_2, L_2\rangle$, then all rules occurring in $\mathcal{A}_1$ also occur in $\mathcal{A}_2$, hence are also verified by all examples of $\langle \mathcal{A}_2, L_2\rangle$. □

The converse of Prop. 12 does not hold. If $\langle \mathcal{A}_2, L_2\rangle^+ \subseteq \langle \mathcal{A}_1, L_1\rangle^+$, then all heads of rules of $\mathcal{A}_1$ must occur in the facts or as heads of rules of $\mathcal{A}_2$, but bodies of rules in $\mathcal{A}_1$ may be different from those in $\mathcal{A}_2$.

4 Priorities based on plausibility

To decide whether an argument not only attacks but also *defeats* another, we need a priority relation among arguments to provide a proper criterion. Here, we will base this priority relation on degrees of plausibility that can be assigned to examples of arguments by ordinal conditional functions. We consider an argument to be as plausible as its most plausible examples. Moreover, the plausibilities of its counterexamples represent the degree to which it can be challenged.

Definition 13 (κ-values of arguments) *Let κ be an ordinal conditional function on Ω, let $\langle \mathcal{A}, L\rangle$ be an argument. Then $\kappa^+(\langle \mathcal{A}, L\rangle) = \min\{\kappa(\omega) \mid \omega \in \langle \mathcal{A}, L\rangle^+\}$, and $\kappa^-(\langle \mathcal{A}, L\rangle) = \min\{\kappa(\omega) \mid \omega \in \langle \mathcal{A}, L\rangle^-\}$.*

When comparing two arguments, the argument the examples of which are more plausible prevails.

Definition 14 (κ-preference) *Let κ be an ordinal conditional function on Ω, let $\langle \mathcal{A}_1, L_1\rangle$, $\langle \mathcal{A}_2, L_2\rangle$ be two arguments. Then $\langle \mathcal{A}_1, L_1\rangle \succeq^\kappa \langle \mathcal{A}_2, L_2\rangle$ iff $\kappa(\mathcal{A}_1{}^+) \leq \kappa(\mathcal{A}_2{}^+)$.*

By using κ-preference as priority relation between DeLP-arguments, an easy criterion for warrant can be derived.

Proposition 15 *Let $\langle \mathcal{A}, L\rangle$ be an argument. If $\kappa^+(\langle \mathcal{A}, L\rangle) < \kappa^-(\langle \mathcal{A}, L\rangle)$, then $\langle \mathcal{A}, L\rangle$ is undefeated and hence warranted.*

Proof: Let $\kappa^+(\langle \mathcal{A}, L\rangle) < \kappa^-(\langle \mathcal{A}, L\rangle)$. Assume there is an argument $\langle \mathcal{A}_1, L_1\rangle$ that (attacks and) defeats $\langle \mathcal{A}, L\rangle$. Then, according to proposition 9 and definition 14, there must be an example $\omega_1 \in \langle \mathcal{A}_1, L_1\rangle^+$ that is a counterexample to $\langle \mathcal{A}, L\rangle$ such that $\kappa(\omega_1) \leq \kappa^+(\langle \mathcal{A}, L\rangle)$. But this contradicts the presupposition $\kappa^+(\langle \mathcal{A}, L\rangle) < \kappa^-(\langle \mathcal{A}, L\rangle)$. $\Box$

Therefore, making use of arbitrary κ-functions allows for a straightforward and intuitive semantical specification of priority between DeLP-arguments. In the following section, we will focus on a specific κ-function – namely system Z – to develop a more syntactical priority relation.

5 Using system Z preference for DeLP

In the original DeLP-framework (García and Simari 2004), *generalized specificity* was used to define a priority relation between arguments, taking the derivation behaviour of arguments into account. In system Z ((Goldszmidt and Pearl 1996) and section 2), a tolerance partitioning of the set of default rules is construed that is also based on a notion of specificity, here applied to single conditionals, not arguments. An idea already mentioned in (García and Simari 2004) will be elaborated in the following: From a priority relation between rules, a corresponding relation between arguments is derived and used for the evaluation of arguments. In this paper, we make use of the Z-ordering on the conditionals in Δ for this purpose and will be able to sharpen proposition 15.

Since we presupposed Δ to be consistent, there is a tolerance partitioning $\Delta = \Delta_0 \cup \ldots \cup \Delta_k$, i.e. the one used by system Z (cf. section 2). Then, we can measure the plausibility of an argument $\langle \mathcal{A}, L\rangle$ by the plausibility of its most plausible rule by setting

$$Z(\langle \mathcal{A}, L\rangle) = \min_{\delta \in \mathcal{A}} Z(\delta).$$

$$\begin{array}{cc} \langle \mathcal{A}_1, f\rangle & U \\ \uparrow & \\ \langle \mathcal{A}_2, \overline{f}\rangle & D \\ \uparrow & \\ \langle \mathcal{A}_3, f\rangle & U \end{array}$$

Figure 1: The dialectical tree for example 17

Note that in fact, the most plausible rules are the weakest links in the chain supporting the conclusion of an argument. Now, arguments can be compared by their Z-values:

Definition 16 (Z-preference) *Argument $\langle \mathcal{A}_1, L_1\rangle$ is* Z-preferred *over argument $\langle \mathcal{A}_2, L_2\rangle$, $\langle \mathcal{A}_1, L_1\rangle \succ^Z \langle \mathcal{A}_2, L_2\rangle$, iff $Z(\langle \mathcal{A}_1, L_1\rangle) > Z(\langle \mathcal{A}_2, L_2\rangle)$. $\langle \mathcal{A}_1, L_1\rangle$ and $\langle \mathcal{A}_2, L_2\rangle$ are* equally Z-preferred*, $\langle \mathcal{A}_1, L_1\rangle \approx^Z \langle \mathcal{A}_2, L_2\rangle$, iff $Z(\langle \mathcal{A}_1, L_1\rangle) = Z(\langle \mathcal{A}_2, L_2\rangle)$.*

To check if a literal is warranted by an argument we have to build up the dialectical tree for this argument from acceptable argumentation lines. Using Z-preference, an argumentation line $[\langle \mathcal{A}_0, L_0\rangle, \langle \mathcal{A}_1, L_1\rangle, \langle \mathcal{A}_2, L_2\rangle, \cdots]$ is acceptable iff all supporting arguments $[\langle \mathcal{A}_0, L_0\rangle, \langle \mathcal{A}_2, L_2\rangle, \langle \mathcal{A}_4, L_4\rangle, \cdots]$, respectively interfering arguments $[\langle \mathcal{A}_1, L_1\rangle, \langle \mathcal{A}_3, L_3\rangle, \langle \mathcal{A}_5, L_5\rangle, \cdots]$ are concordant, and each blocking defeater is followed by a proper defeater, i.e. $\langle \mathcal{A}_{m+1}, L_{m+1}\rangle \succeq^Z \langle \mathcal{A}_m, L_m\rangle$ and $\langle \mathcal{A}_{m+1}, L_{m+1}\rangle \succ^Z \langle \mathcal{A}_{m-1}, L_{m-1}\rangle$ must hold.

We will check which literals can be warranted in our penguin & chicken example 1.

Example 17 We consider the *de.l.p.* $\mathcal{P}_1 = (\{cs\}, \Delta)$, with propositionals and conditionals as specified in example 1. The Z-values of the arguments $\langle \mathcal{A}_1, f\rangle, \langle \mathcal{A}_2, \overline{f}\rangle$ and $\langle \mathcal{A}_3, f\rangle$ as given in example 2 are computed as $Z(\langle \mathcal{A}_1, f\rangle) = 0, Z(\langle \mathcal{A}_2, \overline{f}\rangle) = 1, Z(\langle \mathcal{A}_3, f\rangle) = 2$, so $\langle \mathcal{A}_2, \overline{f}\rangle$ properly defeats $\langle \mathcal{A}_1, f\rangle$, and $\langle \mathcal{A}_3, f\rangle$ properly defeats $\langle \mathcal{A}_2, \overline{f}\rangle$. Hence, $\langle \mathcal{A}_1, f\rangle$ is finally undefeated in the dialectical tree which is shown in figure 1, so f is warranted.

In the example above, the warrant procedure proved to be in accordance with system Z. In the next example, argumentation is shown to do even better than system Z.

Example 18 We work again in the setting of example 1, now considering the defeasible logic program $\mathcal{P}_2 = (\{p\}, \Delta)$. We will check if w is warranted, i.e. if *penguins have wings* can be proved in our argumentation framework. The only argument that can be built to connect p and w is $\langle \{(b|p), (w|b)\}, w\rangle$, which is not attacked at all, so, in particular, is undefeated. Hence, w can be warranted.

On the other hand, in system Z, we have $\kappa_z(pw) = \kappa_z(p\overline{w}) = 1$, since each of the models to be considered falsifies at least one of $(b|p), (f|b)$, or $(\overline{f}|p)$, the most plausible of which – $(f|b)$ – is in Δ_0, just as $(w|b)$. So, the status of the query w can not be determined by system Z. This effect has become known as the *drowning effect* (see, e.g. (Goldszmidt and Pearl 1996)).

We will now prove some general results on arguments and their examples. The following lemma is straightforward, taking into regard that subarguments are actually subsets of their superarguments.

Lemma 19 *Let $\langle \mathcal{A}_1, L_1 \rangle$ be a subargument of $\langle \mathcal{A}_2, L_2 \rangle$. Then $Z(\mathcal{A}_1) \geq Z(\mathcal{A}_2)$.*

Next, we elaborate a relationship between the Z-value of an argument and its counterexamples.

Lemma 20 *Let $\langle \mathcal{A}, L \rangle$ be an argument. Then for all counterexamples $\omega \in \langle \mathcal{A}, L \rangle^-$, $\kappa_z(\omega) \geq Z(\mathcal{A}) + 1$.*

Proof: Each $\omega \in \langle \mathcal{A}, L \rangle^-$ falsifies at least one rule δ in $\mathcal{A}$ with rank $Z(\delta) \geq Z(\mathcal{A})$, so $\kappa_z(\omega) \geq Z(\mathcal{A}) + 1$. □

Finally, we show how the warrant of an argument can be ensured by "good" examples in this framework.

Proposition 21 *Let $\langle \mathcal{A}, L \rangle$ be an argument. If $\kappa_z^+(\langle \mathcal{A}, L \rangle) \leq Z(\langle \mathcal{A}, L \rangle)$, then $\langle \mathcal{A}, L \rangle$ is undefeated, hence warranted.*

Proof: Assume there is $\langle \mathcal{A}_1, L_1 \rangle$ that defeats $\langle \mathcal{A}, L \rangle$ at subargument $\langle \mathcal{A}', L' \rangle$. Then $Z(\langle \mathcal{A}_1, L_1 \rangle) \geq Z(\langle \mathcal{A}', L' \rangle) \geq Z(\langle \mathcal{A}, L \rangle)$ and $L' = \overline{L_1}$. Each example ω of $\langle \mathcal{A}', L' \rangle$ falsifies at least one rule in $\langle \mathcal{A}_1, L_1 \rangle$ so $\kappa^z(\omega) \geq Z(\langle \mathcal{A}_1, L_1 \rangle) + 1 \geq Z(\langle \mathcal{A}, L \rangle) + 1$ for all $\omega \in \langle \mathcal{A}', L' \rangle^+$. Since $\langle \mathcal{A}, L \rangle^+ \subseteq \langle \mathcal{A}', L' \rangle^+$, by Proposition 12, and $\langle \mathcal{A}, L \rangle^+ = \mathit{Mod}(\Phi \wedge \bigwedge_{\delta \in \mathcal{A}} \mathit{head}(\delta))$, by Proposition 6, this implies $\kappa^z(\Phi \wedge \bigwedge_{\delta \in \mathcal{A}} \mathit{head}(\delta)) > Z(\langle \mathcal{A}, L \rangle)$, a contradiction. □

The following corollary is a straightforward consequence.

Corollary 22 *Let $\langle \mathcal{A}, L \rangle$ be an argument. If there is an example ω of $\langle \mathcal{A}, L \rangle$ such that $\kappa_z(\omega) \leq Z(\langle \mathcal{A}, L \rangle)$, then $\langle \mathcal{A}, L \rangle$ is undefeated, hence warranted.*

Example 23 We apply these last results to our penguin & chicken example. For $\langle \mathcal{A}_3, f \rangle$, we compute $\kappa_z^+(\langle \mathcal{A}_3, f \rangle) = \kappa(csf) = 2$ and $Z(\langle \mathcal{A}_3, f \rangle) = 2$, so from proposition 21, we can conclude that f is warranted by $\langle \mathcal{A}_3, f \rangle$, identifying $\langle \mathcal{A}_3, f \rangle$ immediately as a "good" argument for f. A "good" example for f in the sense of corollary 22 can be given by $\omega_1 = csb\overline{p}fw$. Note that $\langle \mathcal{A}_1, f \rangle$ can not be established as a "good" argument for f by proposition 21, since $\kappa_z^+(\langle \mathcal{A}_1, f \rangle) = \kappa(csbf) = 2$, whereas $Z(\langle \mathcal{A}_1, f \rangle) = 0$.

From the last part of example 23, it is obvious that proposition 21 states a sufficient but not necessary condition for warrant in the DeLP-framework.

6 Conclusion

The novel combination of defeasible argumentation with methods from default and conditional reasoning elaborated in this paper was shown to be fruitful in several respects: First, a possible worlds semantics widely used in default reasoning provides a novel semantics for DeLP arguments by taking possible worlds as examples. This allows for attaching a notion of plausibility to arguments in a straightforward way, thus being able to associate a meta-logical priority criterion to the DeLP-framework. Finally, we considered the particular system Z approach to plausibility and proved that in some cases, the complex dialectical evaluation procedure of DeLP can be circumvented by directly comparing the Z-values of examples to the Z-values of the arguments. In this way, each of the two frameworks enriches the other one by new methods and insights. It must be emphasized that DeLP features a fully dialectical frame for argumentation more expressive than other approaches to argumentative inference. In (Benferhat, Dubois, and Prade 1993), argumentative inference is based only on comparing degrees of possibility so that attack there is the same as defeat whereas DeLP also uses the notion of defense to warrant arguments.

Further work might continue in different directions: For example, the Z-value of arguments here was based on the minimum of the Z-degrees of the involved rules, taking the most general (i.e. weakest) rules as a reference point. Other choices are possible, one might go for the most specific rules here, or compare the conflicting rules directly. The aim to be kept in mind would be to specify the results of DeLP's dialectical inference procedure most appropriately in a declarative way. On the other hand, with a focus on default reasoning, one might investigate in more depth in which cases system Z inferences can be realized as results of argumentation processes. Either way, although default reasoning and argumentation are both techniques usable for commonsense reasoning, our impression is that these fields rather complement one another in many fruitful ways than compete with each other.

References

Benferhat, S.; Dubois, D.; and Prade, H. 1993. Argumentative inference in uncertain and inconsistent knowledge bases. In *Proc. 9th Annual Conference on Uncertainty in Artificial Intelligence, UAI-93*, 411–419.

Besnard, P., and Hunter, A. 2008. *Elements of Argumentation*. MIT Press.

DeFinetti, B. 1974. *Theory of Probability*, volume 1,2. New York: John Wiley and Sons.

Dung, P. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77:321–357.

García, A. J., and Simari, G. R. 2004. Defeasible logic programming: An argumentative approach. *Theory and Practice of Logic Programming*, 4(1):95–138.

Goldszmidt, M., and Pearl, J. 1996. Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artificial Intelligence* 84:57–112.

Spohn, W. 1988. Ordinal conditional functions: a dynamic theory of epistemic states. In Harper, W., and Skyrms, B., eds., *Causation in Decision, Belief Change, and Statistics, II*. Kluwer Academic Publishers. 105–134.

Thimm, M., and Kern-Isberner, G. 2008. On the relationship of defeasible argumentation and answer set programming. In Besnard, P.; Doutre, S.; and Hunter, A., eds., *Proceedings of the 2nd International Conference on Computational Models of Argument COMMA'08*, 393–404. IOS Press.