

On the Relation Between Approximation Fixpoint Theory and Justification Theory

Simon Marynissen^{1,2}, Bart Bogaerts² and Marc Denecker¹

¹KU Leuven

²Vrije Universiteit Brussel

simon.marynissen@kuleuven.be, bart.bogaerts@vub.be, marc.denecker@kuleuven.be

Abstract

Approximation Fixpoint Theory (AFT) and Justification Theory (JT) are two frameworks to unify logical formalisms. AFT studies semantics in terms of fixpoints of lattice operators, and JT in terms of so-called justifications, which are explanations of why certain facts do or do not hold in a model. While the approaches differ, the frameworks were designed with similar goals in mind, namely to study the different semantics that arise in (mainly) non-monotonic logics. The first contribution of our current paper is to provide a formal link between the two frameworks. To be precise, we show that every justification frame induces an approximator and that this mapping from JT to AFT preserves all major semantics. The second contribution exploits this correspondence to extend JT with a novel class of semantics, namely *ultimate semantics*: we formally show that ultimate semantics can be obtained in JT by a syntactic transformation on the justification frame, essentially performing some sort of resolution on the rules.

1 Introduction

In this framework, we are concerned with two theories developed with similar intentions, namely to unify semantics of (mostly non-monotonic) logics, namely Approximation Fixpoint Theory (AFT) and Justification Theory (JT).

1.1 Approximation Fixpoint Theory

In the 1980s and 90s, the area of non-monotonic reasoning (NMR) saw fierce debates about formal semantics. In several subareas, researchers sought to formalise common-sense intuitions about knowledge of introspective agents. In these areas, appeals to similar intuitions were made, resulting in the development of similar mathematical concepts. Despite the obvious similarity, the precise relation between these concepts remained elusive. AFT was founded in the early 2000s by Denecker, Marek and Truszczyński [2000] as a way of unifying semantics that emerged in these different subareas. The main contribution of AFT was to demonstrate that, by moving to an algebraic setting, the common principles behind these concepts can be isolated and studied in a general way.

This breakthrough allowed results that were achieved in the context of one of these languages to be easily transferred to another. In the early stages, AFT was applied to default logic, auto-epistemic logic, and logic programming [Denecker *et al.*, 2000; Denecker *et al.*, 2003]. In recent years also applications in various other domains have emerged [Strass, 2013; Bi *et al.*, 2014; Charalambidis *et al.*, 2018; Bogaerts and Cruz-Filipe, 2018].

The foundations of AFT lie in Tarski's fixpoint theory of monotone operators on a complete lattice [Tarski, 1955]. AFT demonstrates that by moving from the original lattice L to the bilattice L^2 , Tarski's theory can be generalised into a fixpoint theory for arbitrary (i.e., also non-monotone) operators. Crucially, all that is required to apply AFT to a formalism and obtain several semantics is to define an appropriate approximating operator $L^2 \rightarrow L^2$ on the bilattice; the algebraic theory of AFT then takes care of the rest. For instance, to characterise the major logic programming semantics using AFT, it suffices to define Fitting's four-valued immediate consequence operator [Fitting, 2002]. The (partial) stable fixpoints of that operator (as defined by AFT) are exactly the partial stable models of the original program; the well-founded fixpoint of the operator is the well-founded model of the program, etc.

1.2 Justification Theory

Building on an old semantical framework for (abductive) logic programming [Denecker and De Schreye, 1993], Denecker *et al.* [2015] defined an abstract theory of *justifications* suitable for describing the semantics of a range of logics in knowledge representation, computational and mathematical logic, including logic programs, argumentation frameworks and nested least and greatest fixpoint definitions. Justifications provide a refined way of describing the semantics of a logic: they not only define whether an interpretation is a model (under a suitable semantics) of a theory, but also *why*.

Justifications — albeit not always in the exact formal form as described by Denecker *et al.* [2015] — have appeared in different ways in different areas. The stable semantics for logic programs was defined in terms of justifications [Fages, 1990; Schulz and Toni, 2013]. Moreover, an algebra for combining justifications (for logic programs) was defined by Cabalar *et al.* [2014]; and justifications are underlying provenance systems in databases [Damásio *et al.*, 2013].

Next to these theoretic benefits, justifications are also used in implementations of answer set solvers (they form the basis of the so-called source-pointer approach in the unfounded set algorithm [Gebser *et al.*, 2009], and turned out to be key in analyzing conflicts in the context of lazy grounding [Bogaerts and Weinzierl, 2018]), as well as to improve parity game solvers [Lapauw *et al.*, 2020].

1.3 Correspondence

The two described frameworks were designed with similar intentions in mind, namely to unify different (mainly non-monotonic) logics. One major difference between them is that JT is defined logically while AFT is defined purely algebraically. This makes justification frameworks less abstract and easier to grasp, but also in a certain sense less general. On the other hand, Denecker *et al.* [2015] defined a notion of nesting, which seems promising to integrate the semantics of nested least and nested greatest fixpoint definitions.

Despite the differences, certain correspondences between the theories show up: several definitions in justification frameworks seem to have an algebraical counterpart in AFT. This is evident from the fact that many results on justifications are formulated in terms of fixpoints of a so-called derivation operator that happens, for the case of logic programming, to coincide with (Fitting’s three-valued version of) the immediate consequence operator for logic programs. Of course, now the question naturally arises whether this correspondence can be made formal, i.e., whether it can formally be shown that semantics induced by JT will always coincide with their equally-named counterpart in AFT. If the answer is positive, this will allow us to translate results between the two theories. Formalising this correspondence is the key contribution of the current paper.

1.4 Contributions

Our contributions can be summarised as follows:

- In Section 3, we provide some novel results for JT. While the main purpose of these results is to support the theorems of Section 4, they also directly advance the state of JT. In this section, we show among others how different semantics induced by JT relate, and we resolve a discrepancy that exists between different definitions of so-called *stable and supported branch evaluations* in prior work. We formally prove that the different circulating definitions of these branch evaluations indeed induce the same semantics.
- In Section 4, we turn our attention to the key contribution of the paper, namely embedding JT in AFT. To do this, we proceed as follows. First, we show that under minor restrictions, each justification frame (intuitively, this is a set of rules that describe when a positive or negative fact is true), can be transformed into an approximator. Next, we show that for each of the most common *branch evaluations* (these are mathematical structures that are used to associate semantics to a justification frame), the induced semantics by JT is the same as the equally-named semantics on the AFT

side. Establishing this result is of particular importance for the future development of JT, since this result immediately makes a large body of theoretical results developed in the context of AFT readily available for JT, as well as all its future application domains, including results on stratification [Vennekens *et al.*, 2006; Bogaerts and Cruz-Filipe, 2021], predicate introduction [Vennekens *et al.*, 2007], and knowledge compilation [Bogaerts and Van den Broeck, 2015]. On the other hand, from the context of AFT, the embedding of JT can serve as inspiration for developing more general algebraic explanation mechanisms.

- To illustrate how this connection can be exploited for further exploiting the theory of justifications, we turn our attention to *ultimate semantics*. In the context of AFT, Denecker and his coauthors have realised that a single operator can have multiple approximators and that the choice of approximator influences the induced semantics. They also showed that — when staying in the realm of consistent AFT — every operator induces a most precise approximator, and called this *the ultimate approximator* [Denecker *et al.*, 2004]. In Section 5, we transfer this idea to JT. We show there that by means of a simple transformation¹ on the justification frame, we can obtain ultimate semantics. Importantly, since this transformation is defined independent of the branch evaluation at hand, ultimate semantics are not just induced for the semantics that have a counterpart in AFT, but for all conceivable current and future branch evaluations as well.

2 Preliminaries: Justification Theory

In this section, we use the formalisation of JT as done by Marynissen *et al.* [2020]. Truth values are denoted **t** (true), **f** (false) and **u** (unknown); we write \mathcal{L} for $\{\mathbf{t}, \mathbf{f}, \mathbf{u}\}$. We make use of two orders on \mathcal{L} , the *truth order* $\mathbf{f} \leq_t \mathbf{u} \leq_t \mathbf{t}$ and the *precision order* $\mathbf{u} \leq_p \mathbf{f}, \mathbf{t}$. JT starts with a set \mathcal{F} , referred to as a *fact space*, such that $\mathcal{L} \subseteq \mathcal{F}$; the elements of \mathcal{F} are called *facts*. We assume that \mathcal{F} is equipped with an involution $\sim: \mathcal{F} \rightarrow \mathcal{F}$ (i.e., a bijection that is its own inverse) such that $\sim \mathbf{t} = \mathbf{f}$, $\sim \mathbf{u} = \mathbf{u}$, and $\sim x \neq x$ for all $x \neq \mathbf{u}$. Moreover, we assume that $\mathcal{F} \setminus \mathcal{L}$ is partitioned into two disjoint sets \mathcal{F}_+ and \mathcal{F}_- such that $x \in \mathcal{F}_+$ if and only if $\sim x \in \mathcal{F}_-$ for all $x \in \mathcal{F} \setminus \mathcal{L}$. Elements of \mathcal{F}_+ are called *positive* and elements of \mathcal{F}_- are called *negative* facts. An example of a fact space is the set of literals over a propositional vocabulary Σ extended with \mathcal{L} where \sim maps a literal to its negation. For any set A we define $\sim A$ to be the set of elements of the form $\sim a$ for $a \in A$. We distinguish two types of facts: *defined* and *open* facts. The former are accompanied by a set of rules that determine their truth value. The truth value of the latter is not governed by the rule system but comes from an external source or is fixed (as is the case for logical facts).

Definition 1. A *justification frame* \mathcal{JF} is a tuple $\langle \mathcal{F}, \mathcal{F}_d, R \rangle$ such that

¹Essentially, this transformation performs some sort of case splitting.

- \mathcal{F}_d is a subset of \mathcal{F} closed under \sim , i.e., $\sim\mathcal{F}_d = \mathcal{F}_d$; facts in \mathcal{F}_d are called *defined*;
- no logical fact is defined: $\mathcal{L} \cap \mathcal{F}_d = \emptyset$;
- $R \subseteq \mathcal{F}_d \times 2^{\mathcal{F} \setminus \emptyset}$;
- for each $x \in \mathcal{F}_d$ there is at least one element $(x, A) \in R$.

The set \mathcal{F}_o of *open* facts is equal to $\mathcal{F} \setminus \mathcal{F}_d$. An element $(x, A) \in R$ is called a *rule* with *head* x and *body* (or *case*) A . The set of cases of x in \mathcal{JF} is denoted $\mathcal{JF}(x)$. Rules $(x, A) \in R$ are denoted as $x \leftarrow A$ and if $A = \{y_1, \dots, y_n\}$, we often write $x \leftarrow y_1, \dots, y_n$.

Logic programming rules can easily be transferred to rules in a justification frame. However, in logic programming, only rules for positive facts are given; never for negative facts. Hence, in order to apply JT to logic programming, a mechanism for deriving rules for negative literals is needed. For this, a technique called *complementation* was invented [Dencker *et al.*, 2015]; it is a generic mechanism that allows turning a set of rules for x into a set of rules for $\sim x$. To define complementation, we first define *selection functions* for x . A selection function for x is a mapping $s: \mathcal{JF}(x) \rightarrow \mathcal{F}$ such that $s(A) \in A$ for all rules of the form $x \leftarrow A$. Intuitively, a selection function chooses an element from the body of each rule of x . For a selection function s , the set $\{s(A) \mid A \in \mathcal{JF}(x)\}$ is denoted by $\text{Im}(s)$.

Definition 2. For a set of rules R , we define R^* to be the set of rules of the form $\sim x \leftarrow \sim \text{Im}(s)$ for $x \in \mathcal{F}_d$ that has rules in R and s a selection function for x . The *complementation* of \mathcal{JF} is defined as $\langle \mathcal{F}, \mathcal{F}_d, R \cup R^* \rangle$. A justification frame \mathcal{JF} is *complementary* if it is fixed under complementation, i.e., $R \cup R^* = R$.


Example 1. If $R = \{x \leftarrow a, b; x \leftarrow c\}$, then $R^* = \{\sim x \leftarrow \sim a, \sim c; \sim x \leftarrow \sim b, \sim c\}$.

Definition 3. A *directed graph* is a pair (N, E) where N is a set of nodes and $E \subseteq N \times N$ is the set of edges. An *internal* node is a node with outgoing edges. A *leaf* is a non-internal node.

Definition 4. Let $\mathcal{JF} = \langle \mathcal{F}, \mathcal{F}_d, R \rangle$ be a justification frame. A *justification* J in \mathcal{JF} is a directed graph (N, E) such that for every internal node $n \in N$ it holds that $n \leftarrow \{m \mid (n, m) \in E\} \in R$;

A justification is *locally complete* if it has no leaves in \mathcal{F}_d . We write $\mathfrak{J}(x)$ to denote the set of locally complete justifications that have an internal node x .

Example 2. Take $\mathcal{F}_d = \{x, \sim x, y, \sim y\}$, $\mathcal{F}_o = \{a, \sim a, b, \sim b\} \cup \mathcal{L}$, and R the complementation of $\{x \leftarrow y, a; y \leftarrow y, b\}$, then



is a locally complete justification in $\langle \mathcal{F}, \mathcal{F}_d, R \rangle$ because a and b are open facts.

Definition 5. Let \mathcal{JF} be a justification frame. A *\mathcal{JF} -branch* is either an infinite sequence in \mathcal{F}_d or a finite non-empty sequence in \mathcal{F}_d followed by an element in \mathcal{F}_o . For a justification J in \mathcal{JF} , a *J -branch* starting from $x \in \mathcal{F}_d$ is a path in J starting from x that is either infinite or ends in a leaf of J . We write $B_J(x)$ to denote the set of J -branches starting from x .

Not all J -branches are \mathcal{JF} -branches since they can end in a defined fact. However, if J is locally complete, any J -branch is also a \mathcal{JF} -branch. We denote a branch \mathbf{b} as $\mathbf{b}: x_0 \rightarrow x_1 \rightarrow \dots$ and define $\sim \mathbf{b}$ as $\sim x_0 \rightarrow \sim x_1 \rightarrow \dots$. A *tail* of a branch \mathbf{b} is a branch $x_i \rightarrow x_{i+1} \rightarrow \dots$ for some $i \geq 0$.

Definition 6. A *branch evaluation* \mathcal{B} is a mapping that maps any \mathcal{JF} -branch to an element in \mathcal{F} for all justification frames \mathcal{JF} . A branch evaluation \mathcal{B} is *consistent* if $\mathcal{B}(\sim \mathbf{b}) = \sim \mathcal{B}(\mathbf{b})$ for any branch \mathbf{b} . A justification frame \mathcal{JF} together with a branch evaluation \mathcal{B} form a *justification system* \mathcal{JS} , which is presented as a quadruple $\langle \mathcal{F}, \mathcal{F}_d, R, \mathcal{B} \rangle$.

The main branch evaluations we are interested in are given below:

Definition 7. The *supported* branch evaluation \mathcal{B}_{sp} maps $x_0 \rightarrow x_1 \rightarrow \dots$ to x_1 . The *Kripke-Kleene* branch evaluation \mathcal{B}_{KK} maps finite branches to their last element and infinite branches to \mathbf{u} . The *well-founded* branch evaluation \mathcal{B}_{wf} maps finite branches to their last element. It maps infinite branches to \mathbf{t} if they have a negative tail, to \mathbf{f} if they have a positive tail and to \mathbf{u} otherwise. The *stable* branch evaluation \mathcal{B}_{st} maps a branch $x_0 \rightarrow x_1 \rightarrow \dots$ to the first element that has a different sign than x_0 if it exists; otherwise \mathbf{b} is mapped to $\mathcal{B}_{\text{wf}}(\mathbf{b})$.

Definition 8. A (*three-valued*) *interpretation* of \mathcal{F} is a function $\mathcal{I}: \mathcal{F} \rightarrow \mathcal{L}$ such that $\mathcal{I}(\sim x) = \sim \mathcal{I}(x)$ for all $x \in \mathcal{F}$ and $\mathcal{I}(\ell) = \ell$ for all $\ell \in \mathcal{L}$.

We will assume that the interpretation of open facts is fixed; hence any two interpretations coincide on open facts.

Definition 9. Let $\mathcal{JS} = \langle \mathcal{F}, \mathcal{F}_d, R, \mathcal{B} \rangle$ be a justification system, \mathcal{I} an interpretation of \mathcal{F} , and J a locally complete justification in \mathcal{JS} . Let $x \in \mathcal{F}_d$ be a node in J . The *value* of $x \in \mathcal{F}_d$ by J under \mathcal{I} is defined as $\text{val}_{\mathcal{B}}(J, x, \mathcal{I}) = \min_{\mathbf{b} \in B_J(x)} \mathcal{I}(\mathcal{B}(\mathbf{b}))$, where \min is with respect to \leq_t .

The *supported value* of $x \in \mathcal{F}$ in \mathcal{JS} under \mathcal{I} is defined as $\text{SV}_{\mathcal{JS}}(x, \mathcal{I}) = \max_{J \in \mathfrak{J}(x)} \text{val}_{\mathcal{B}}(J, x, \mathcal{I})$ for $x \in \mathcal{F}_d$ and $\text{SV}(x, \mathcal{I}) = \mathcal{I}(x)$ for $x \in \mathcal{F}_o$. For any \mathcal{F} -interpretation \mathcal{I} , $\mathcal{S}_{\mathcal{JS}}(\mathcal{I})$ is the mapping $\mathcal{F} \rightarrow \mathcal{L}: x \mapsto \text{SV}_{\mathcal{JS}}(x, \mathcal{I})$. The function $\mathcal{S}_{\mathcal{JS}}$ is called the *support operator*. If \mathcal{JS} consists of \mathcal{JF} and \mathcal{B} , then we write $\mathcal{S}_{\mathcal{JF}}^{\mathcal{B}}$ for $\mathcal{S}_{\mathcal{JS}}$. If \mathcal{JF} is clear from context, we write $\text{SV}_{\mathcal{B}}$ for $\text{SV}_{\mathcal{JS}}$.


Models under justification semantics are determined by the supported value.

Definition 10. Let \mathcal{JF} be justification frame and \mathcal{B} a branch evaluation. An \mathcal{F} -interpretation \mathcal{I} is a *\mathcal{B} -model* of \mathcal{JF} if for all $x \in \mathcal{F}$, $\text{SV}_{\mathcal{JF}}^{\mathcal{B}}(x, \mathcal{I}) = \mathcal{I}(x)$, i.e., \mathcal{I} is a fixpoint of $\mathcal{S}_{\mathcal{JF}}^{\mathcal{B}}$.

A \mathcal{B}_{sp} , \mathcal{B}_{KK} , \mathcal{B}_{st} , or \mathcal{B}_{wf} -model is called a supported, Kripke-Kleene, stable, or well-founded model.

Example 3. Let $\mathcal{F} = \{p, \sim p, q, \sim q\} \cup \mathcal{L}$ and take R to be the complementation of $\{p \leftarrow p; p \leftarrow \sim q; q \leftarrow q\}$, i.e. adding

the rules $\sim p \leftarrow \sim p, q$ and $\sim q \leftarrow \sim q$. There are exactly two locally complete justifications with p as an internal node:



Under \mathcal{B}_{wf} , the left justification has a value \mathbf{f} for p , while the right justification has a value \mathbf{t} for p . Since these justifications are the only locally complete ones containing p as an internal node we have that $\text{SV}_{\mathcal{B}_{\text{wf}}}(p, \mathcal{I}) = \mathbf{t}$ for all interpretations \mathcal{I} of \mathcal{F} . This shows us that the unique \mathcal{B}_{wf} -model is the interpretation mapping p to \mathbf{t} and q to \mathbf{f} .

3 Tying Up Loose Ends

In this section, we prove some results about JT that will be needed for developing our theory later on. These results resolve several issues that were left open in prior work, but turn out to be crucial for studying the relationship with AFT.

3.1 Pasting Justifications

Our first result is essentially a pasting result. What it states is that we can, for all branch evaluations of interest to the current paper, build a single justification that explains the value of all facts. In other words, it provides a means of gluing justifications for different facts together. The first theorem has already been proven in a slightly different context by Marynissen *et al.* [2018].

Theorem 1. Take $\mathcal{B} \in \{\mathcal{B}_{\text{sp}}, \mathcal{B}_{\text{KK}}, \mathcal{B}_{\text{st}}, \mathcal{B}_{\text{wf}}\}$. For every interpretation \mathcal{I} , there is a locally complete justification J such that $\text{val}_{\mathcal{B}}(x, J, \mathcal{I}) = \text{SV}_{\mathcal{B}}(x, \mathcal{I})$ for all $x \in \mathcal{F}_d$.

Similarly, this holds for \mathcal{B}'_{sp} and \mathcal{B}'_{st} , but only for models.

Theorem 2. Take $\mathcal{B} \in \{\mathcal{B}'_{\text{sp}}, \mathcal{B}'_{\text{st}}\}$. For every \mathcal{B} -model \mathcal{I} , there is a locally complete justification J such that $\text{val}_{\mathcal{B}}(x, J, \mathcal{I}) = \text{SV}_{\mathcal{B}}(x, \mathcal{I})$ for all $x \in \mathcal{F}_d$.

3.2 Equivalence of Branch Evaluations

Our second result concerns different versions of the stable and supported branch evaluations that circulate in prior work. Marynissen *et al.* [2018; 2020] use the stable and supported branch evaluations as we defined them in Definition 7, while Denecker *et al.* [2015] use the following alternative.

Definition 11. The branch evaluation \mathcal{B}'_{sp} is equal to \mathcal{B}_{sp} on infinite branches and maps finite branches to their last element. The branch evaluation \mathcal{B}'_{st} is equal to \mathcal{B}_{st} except \mathcal{B}'_{st} maps any finite branch to its last element.

Of course this begs the question in which sense these branch evaluations are related. What we show next is that for the purpose of defining models, they are interchangeable.

Definition 12. Two branch evaluations \mathcal{B}_1 and \mathcal{B}_2 are *equivalent* if for all justification frames \mathcal{JF} , the \mathcal{B}_1 -models and the \mathcal{B}_2 -models of \mathcal{JF} coincide.

Our proofs that \mathcal{B}_{sp} and \mathcal{B}'_{sp} , and \mathcal{B}_{st} and \mathcal{B}'_{st} are equivalent, will make use of the following lemma, which intuitively states that to show that an interpretation is a \mathcal{B} -model, it suffices to show that the supported value of each fact is *at least* its value in the interpretation.

Lemma 1. Take $\mathcal{JS} = \langle \mathcal{F}, \mathcal{F}_d, R, \mathcal{B} \rangle$ with \mathcal{B} consistent. Every interpretation \mathcal{I} such that $\text{SV}_{\mathcal{JS}}(x, \mathcal{I}) \geq_t \mathcal{I}(x)$ for all $x \in \mathcal{F}_d$, is a \mathcal{B} -model of \mathcal{JF} .

Proposition 1. The two supported branch evaluations \mathcal{B}_{sp} and \mathcal{B}'_{sp} are equivalent.

Proposition 2. The two stable branch evaluations \mathcal{B}_{st} and \mathcal{B}'_{st} are equivalent.

Sketch of the proofs of Propositions 1 and 2. The difference between the two branch evaluations at hand is that in the one (\mathcal{B}'), finite branches are evaluated with respect to their final element, and the other (\mathcal{B}) with respect to some other element y in the branch (second or first sign switch). Take a \mathcal{B} -model \mathcal{I} and take a justification J as from Theorems 1 or 2. With some care, we can prove that the final element of a J -branch has a larger value than first element in \mathcal{I} under \leq_t . Therefore, \mathcal{I} satisfies the conditions of Lemma 1, proving that \mathcal{I} is a \mathcal{B}' -model. The other direction is proven similarly. \square

3.3 Links between Different Justification Models

Our third set of results is concerned with the relation between different semantics induced by JT. In the context of logic programming, it is well-known that there is a unique well-founded model, what the relation between well-founded and stable model is, etcetera. Several such results will follow immediately by establishing the correspondence with AFT, but some of them will be needed in our proof. They are given explicitly, and sometimes in higher generality, in the current section.

First of all, in logic programming, it is well-known that the well-founded and Kripke-Kleene semantics induce a single model. In JT, we find an analogous result for a broad class of branch evaluations. A branch evaluation \mathcal{B} is called *parametric* if $\mathcal{B}(\mathbf{b}) \in \mathcal{F}_o$ for all \mathcal{JF} -branches and all justification frames \mathcal{JF} . Denecker *et al.* [2015] provided the following result.

Proposition 3. If \mathcal{JF} is a justification frame and \mathcal{B} a parametric branch evaluation, then \mathcal{JF} has a single \mathcal{B} -model.

In this proposition, we of course make use of our earlier assumption that the value of the open facts is fixed. In general, every interpretation of the open facts induces a single model.

Corollary 1. Every justification frame has a unique \mathcal{B}_{KK} -model and a unique \mathcal{B}_{wf} -model.

Proposition 4. The unique \mathcal{B}_{KK} -model is a \mathcal{B}_{sp} -model.

Proposition 5. The well-founded model of \mathcal{JF} is a stable model of \mathcal{JF} .

Proposition 6. Every stable justification model is a supported justification model.

Proof sketches of Proposition 4, 5 and 6. The idea of the proof is that the justification according to Theorem 1 also works for the other branch evaluation. \square

Lemma 2. *Let \mathcal{I} be a stable justification model. If $\text{SV}_{\mathcal{B}_{\text{wf}}}(x, \mathcal{I}) = \mathbf{f}$, then $\mathcal{I}(x) = \mathbf{f}$. If $\text{SV}_{\mathcal{B}_{\text{wf}}}(x, \mathcal{I}) = \mathbf{t}$, then $\mathcal{I}(x) = \mathbf{t}$.*

Proof sketch. Take $x \in \mathcal{F}_d$ with $\mathcal{I}(x) = \mathbf{f}$. By using that \mathcal{I} is a \mathcal{B}_{st} -model, we can prove that every justification J has a branch $\mathbf{b} \in B_J(x)$ such that $\mathcal{I}(\mathcal{B}_{\text{wf}}(\mathbf{b})) \leq_t \mathbf{u}$. This concludes that $\text{SV}_{\mathcal{B}_{\text{wf}}}(x, \mathcal{I}) \neq \mathbf{t}$. \square

Proposition 7. *The well-founded model is the \leq_p -least stable model.*

Proof sketch. Lemma 2 implies that the well-founded model is less precise than any stable model. Then Proposition 5 finishes the proof. \square

4 Embedding JT in AFT

We now turn our attention to the main topic of this paper, namely formally prove the correspondence between JT and AFT. We start with a brief recall of the basic definitions that constitute AFT, next show how to obtain an approximator out of a justification frame, and finally prove that indeed, all major semantics are preserved under this correspondence.

4.1 Preliminaries: AFT

Given a complete lattice $\langle L, \leq \rangle$, AFT [Denecker *et al.*, 2000] uses the *bilattice* $L^2 = L \times L$. We define projection functions as usual: $(x, y)_1 = x$ and $(x, y)_2 = y$. Pairs $(x, y) \in L^2$ are used to approximate elements in the interval $[x, y] = \{z \mid x \leq z \leq y\}$. We call $(x, y) \in L^2$ *consistent* if $x \leq y$, i.e., if $[x, y]$ is not empty. The set of consistent elements is denoted L^c . A pair (x, x) is called *exact* since it approximates only the element x . The *precision order* \leq_p on L^2 is defined as $(x, y) \leq_p (u, v)$ if $x \leq u$ and $y \geq v$. If (u, v) is consistent, this means that $[u, v] \subseteq [x, y]$. If $\langle L, \leq \rangle$ is a complete lattice, then so is $\langle L^2, \leq_p \rangle$. AFT studies fixpoints of operators $O: L \rightarrow L$ through operators approximating O . An operator $A: L^2 \rightarrow L^2$ is an *approximator* of O if it is \leq_p -monotone and has the property that $A(x, x) = (O(x), O(x))$ for all $x \in L$. Approximators are internal in L^c (i.e., map L^c into L^c). We often restrict our attention to *symmetric* approximators: approximators A such that, for all x and y , $A(x, y)_1 = A(y, x)_2$. Denecker *et al.* [2004] showed that the consistent fixpoints of interest of a symmetric approximator are uniquely determined by an approximator's restriction to L^c and hence, that it usually suffices to define approximators on L^c . Such a restriction is called a *consistent approximator*. As mentioned before, AFT studies fixpoints of O using fixpoints of A . The main type of fixpoints that concern us are given here.

- A *partial supported fixpoint* of A is a fixpoint of A .
- The *Kripke-Kleene fixpoint* of A is the \leq_p -least fixpoint of A ; it approximates all fixpoints of A .
- A *partial stable fixpoint* of A is a pair (x, y) such that $x = \text{lfp}(A(\cdot, y)_1)$ and $y = \text{lfp}(A(x, \cdot)_2)$, where $A(\cdot, y)_1$ denotes the function $L \rightarrow L: z \mapsto A(z, y)_1$ and analogously $A(x, \cdot)_2$ stands for $L \rightarrow L: z \mapsto A(x, z)_2$.
- The *well-founded fixpoint* of A is the \leq_p -least partial stable fixpoint of A .

4.2 The Approximator

Let $\mathcal{JF} = \langle \mathcal{F}, \mathcal{F}_d, R \rangle$ be a justification frame, fixed throughout this section. Our first goal is to define, from a given justification frame, an approximator on a suitable lattice. Following the correspondence with how this is done in logic programming, we will take as lattice the set of *exact* interpretations (interpretations that map no facts to \mathbf{u} except for \mathbf{u} itself). It is easy to see that such interpretations correspond directly to subsets of \mathcal{F}_+ . In other words, we will use the lattice $\langle L = 2^{\mathcal{F}_+}, \subseteq \rangle$. Now, the set L^c is isomorphic to the set of three-valued interpretations of \mathcal{F} ; under this isomorphism, a consistent pair $(I, J) \in L^c$ corresponds to the three-valued interpretation \mathcal{I} such that for positive facts $x \in \mathcal{F}_+$, $\mathcal{I}(x) = \mathbf{t}$ if $x \in I$, $\mathcal{I}(x) = \mathbf{f}$ if $x \notin J$, and $\mathcal{I}(x) = \mathbf{u}$ otherwise.

Definition 13. The operator $O_{\mathcal{JF}}: L \rightarrow L$ of \mathcal{JF} maps a subset I of \mathcal{F}_+ to

$$O_{\mathcal{JF}}(I) = \{x \in \mathcal{F}_+ \mid \exists x \leftarrow A \in R: \forall a \in A: (I, I)(a) = \mathbf{t}\}.$$

The *approximator* $A_{\mathcal{JF}}: L^c \rightarrow L^c$ of \mathcal{JF} is defined as follows

$$A_{\mathcal{JF}}(\mathcal{I})_1 = \{x \in \mathcal{F}_+ \mid \exists x \leftarrow A \in R: \forall a \in A: \mathcal{I}(a) = \mathbf{t}\}$$

$$A_{\mathcal{JF}}(\mathcal{I})_2 = \{x \in \mathcal{F}_+ \mid \exists x \leftarrow A \in R: \forall a \in A: \mathcal{I}(a) \geq_t \mathbf{u}\}$$

Proposition 8. *If no rule body in \mathcal{JF} contains \mathbf{u} , then $A_{\mathcal{JF}}$ is a consistent approximator of $O_{\mathcal{JF}}$.*

So far, we are not aware of practical examples with bodies containing \mathbf{u} . From now on, we assume that every justification frame does not have \mathbf{u} in a rule body. It turns out that in case our justification frame behaves well with respect to negation (if it is complementary), the approximator is the same operator as induced by the branch evaluation \mathcal{B}_{sp} .

Lemma 3. *For a complementary justification frame \mathcal{JF} , the function $A_{\mathcal{JF}}$ and the support operator $\mathcal{S}_{\mathcal{JF}}^{\text{Bsp}}$ are equal.*

4.3 Semantic Correspondence

The central result of this section is the following theorem, which essentially states that for all major semantics, the branch evaluation in JT corresponds to the definitions of AFT.

Theorem 3. *Take a complementary justification frame \mathcal{JF} .*

- *The partial supported fixpoints of $A_{\mathcal{JF}}$ are exactly the supported models of \mathcal{JF} .*
- *The Kripke-Kleene fixpoint of $A_{\mathcal{JF}}$ is the unique Kripke-Kleene model of \mathcal{JF} .*
- *The partial stable fixpoints of $A_{\mathcal{JF}}$ are exactly the stable models of \mathcal{JF} .*
- *The well-founded fixpoint of $A_{\mathcal{JF}}$ is the unique well-founded model of \mathcal{JF} .*

These four points are proven independently; the first follows directly from our observation that $A_{\mathcal{JF}}$ and $\mathcal{S}_{\mathcal{JF}}^{\text{Bsp}}$ are in fact the same operator.

Proposition 9. *The partial supported fixpoints of $A_{\mathcal{JF}}$ are exactly the supported models of \mathcal{JF} .*

Given the correspondence between supported semantics, the result for Kripke-Kleene semantics follows quite easily.

Proposition 10. *The Kripke-Kleene fixpoint of $A_{\mathcal{JF}}$ is equal to the unique \mathcal{B}_{KK} -model of \mathcal{JF} .*

Proof sketch. By combining Propositions 4, and 9, we get that the unique \mathcal{B}_{KK} -model (denoted here $\mathcal{I}_{\mathcal{B}_{\text{KK}}}$) is a fixpoint of $A_{\mathcal{JF}}$. All that is left to show is that it is the least precise one. Assume towards contradiction that this is not the case, i.e., that there is a fixpoint \mathcal{I} of $A_{\mathcal{JF}}$ such that $\mathcal{I}_{\mathcal{B}_{\text{KK}}} \not\leq_p \mathcal{I}$. From this we can find an $x \in \mathcal{F}_d$ such that either $\mathcal{I}_{\mathcal{B}_{\text{KK}}}(x) = \mathbf{f}$ and $\mathcal{I}(x) = \mathbf{t}$, or $\mathcal{I}_{\mathcal{B}_{\text{KK}}}(x) = \mathbf{f}$ and $\mathcal{I}(x) = \mathbf{u}$. In both cases, we have $\mathcal{I}_{\mathcal{B}_{\text{KK}}}(x) = \mathbf{f}$; hence every justification J with x as internal node has a finite branch starting with x mapped to an open fact y with $\mathcal{I}_{\mathcal{B}_{\text{KK}}}(y) = \mathbf{f}$. However, since $\mathcal{I}_{\mathcal{B}_{\text{KK}}}$ and \mathcal{I} agree on open facts, this also means that $\text{SV}_{\mathcal{B}'_{\text{sp}}}(x, \mathcal{I}) = \mathbf{f}$, contradicting that \mathcal{I} is a \mathcal{B}'_{sp} -model. \square

The proof of the third point of Theorem 3 is split in two parts, proven separately in the following propositions.

Proposition 11. *Each stable model of \mathcal{JF} is a partial stable fixpoint of $A_{\mathcal{JF}}$.*

Proof sketch. Take a \mathcal{B}_{st} -model $\mathcal{I} = (I_1, I_2)$ of \mathcal{JF} . To show that \mathcal{I} is a partial stable fixpoint, we have two lfp equations to prove. We focus on $\text{lfp}(A_{\mathcal{JF}}(\cdot, I_2)_1) = I_1$. We know from Theorem 1 that a justification J exists that justifies all facts in I_1 . Now, this specific justification induces a dependency order on the facts in I_1 , defined as $y \preceq_J x$ if y is reachable in J from x through positive facts. Using the definition of stable branch evaluation, we can see that this order is well-founded and subsequently prove using well-founded induction on this order that all facts in I_1 must be in $\text{lfp}(A_{\mathcal{JF}}(\cdot, I_2)_1)$. \square

For the other direction, we first need the following lemma.

Lemma 4. *Let \mathcal{I} be a \mathcal{B}_{sp} -model and $x \in \mathcal{F}_+$ with $\text{SV}_{\mathcal{B}_{\text{sp}}}(x, \mathcal{I}) = \mathbf{f}$. It holds that $\text{SV}_{\mathcal{B}_{\text{st}}}(x, \mathcal{I}) = \mathbf{f}$.*

Proposition 12. *Each partial stable fixpoint of $A_{\mathcal{JF}}$ is a stable model of \mathcal{JF} .*

Proof sketch. Let $\mathcal{I} = (I_1, I_2)$ be a partial stable fixpoint of $A_{\mathcal{JF}}$. We prove that $\text{SV}_{\mathcal{B}_{\text{st}}}(x, \mathcal{I}) = \mathcal{I}(x)$ for $x \in \mathcal{F}_+$. By consistency of $\mathcal{S}_{\mathcal{JF}}^{\mathcal{B}_{\text{st}}}$ [Marynissen *et al.*, 2018], this proves that \mathcal{I} is a stable model of \mathcal{JF} . The proof consists of three parts.

1. $\text{SV}_{\mathcal{B}_{\text{st}}}(x, \mathcal{I}) = \mathcal{I}(x) = \mathbf{t}$ for all $x \in I_1$.
2. $\text{SV}_{\mathcal{B}_{\text{st}}}(x, \mathcal{I}) = \mathcal{I}(x) = \mathbf{u}$ for all $x \in I_2 \setminus I_1$.
3. $\text{SV}_{\mathcal{B}_{\text{st}}}(x, \mathcal{I}) = \mathcal{I}(x) = \mathbf{f}$ for all $x \in \mathcal{F}_+ \setminus I_2$.

The first part is proven by constructing a possibly non-locally complete justification without infinite branches, every internal node positive, and for every node y we have that $\mathcal{I}(y) = \mathbf{t}$. The construction is possible because I_1 is the least fixpoint of $A_{\mathcal{JF}}(\cdot, I_2)_1$. The second part is proven similarly. Last part is a consequence of Lemma 4. \square

Example 4. *Let $\mathcal{F} = \{x, \sim x, y, \sim y, z, \sim z\} \cup \mathcal{L}$ and let R be the complementation of $\{x \leftarrow y; y \leftarrow \sim z; z \leftarrow \sim x, \sim y\}$, i.e., it adds the rules $\sim x \leftarrow \sim y; \sim y \leftarrow z; \sim z \leftarrow x$ and $\sim z \leftarrow y$. The approximator $A_{\mathcal{JF}}$ has three partial stable fixpoints: $(\{x, y\}, \{x, y\})$, $(\{z\}, \{z\})$ and $(\emptyset, \{x, y, z\})$.*

Let us take a look at the fixpoint $(\{x, y\}, \{x, y\})$. Since it is a stable fixpoint, we know that $(\{x, y\}, \{x, y\})$ is a least fixpoint of $A_{\mathcal{JF}}(\cdot, \{x, y\})$. This operator is monotone with respect to \subseteq ; hence we can construct the fixpoint by iteratively applying the operator on $(\emptyset, \{x, y\})$. This produces the following sequence.

$$(\emptyset, \{x, y\}) \rightarrow (\{y\}, \{x, y\}) \rightarrow (\{x, y\}, \{x, y\})$$

The first step uses the rule $y \leftarrow \sim z$, while the second step uses the rule $x \leftarrow y$. Combining the two we get the justification $x \rightarrow y \rightarrow \sim z$, which has only true nodes in the model $(\{x, y\}, \{x, y\})$, only positive internal nodes and every defined leaf is negative. This illustrates the first step of the proof of Proposition 12. By extending the found justification, we get a locally complete justification with the same value as the supported value.

The proof of the fourth point of Theorem 3 follows directly from the third point and Proposition 7.

5 Ultimate Semantics for Justification Frames

When applying AFT to new domains, there is not always a clear choice of which approximator to use; the operator on the other hand is often more clear. Denecker *et al.* [2004] studied the space of approximators and observed that consistent approximators can naturally be ordered according to their precision, where more precise approximator also yield more precise results (e.g., if approximator A is more precise than B , then the A -well-founded fixpoint is guaranteed to be more precise than B 's). They also observed that the space of approximators of O has a most precise element, called the *ultimate approximator*, denoted $U(O)$.

In the context of JT, the justification frame uniquely determines the approximator at hand. Still, we show that it is possible to obtain ultimate semantics here as well. To do so, we will develop a method to transform a justification frame \mathcal{JF} into its ultimate frame $U(\mathcal{JF})$. We will then show that the approximator associated to $U(\mathcal{JF})$ is indeed the ultimate approximator of $O_{\mathcal{JF}}$. The result is a generic mechanism to go from any semantics induced by JT (for arbitrary branch evaluations – not just for those that have an AFT counterpart) to an ultimate variant thereof. Our construction is as follows:

Definition 14. Let \mathcal{JF} be a complementary justification frame. Let X be the set rules with a positive head. Let X^* be the least (w.r.t. \subseteq) set containing X that is closed under the addition of rules $x \leftarrow A$ if there is a rule $x \leftarrow B$ with $B \subseteq A$, or if there are rules $x \leftarrow \{y\} \cup A$, $x \leftarrow \{\sim y\} \cup A$. Let Y be the complementation of X . Then $U(\mathcal{JF})$ is defined to be the complementary justification frame $\langle \mathcal{F}, \mathcal{F}_d, Y \rangle$.

Example 5. *Let $\mathcal{F}_d = \{x, \sim x\}$ and $\mathcal{F}_o = \{a, \sim a, b, \sim b\} \cup \mathcal{F}_o$. Take R to be the complementation of $\{x \leftarrow a, x \leftarrow b, \sim x\}$, which adds the following rules*

$$\begin{array}{ll} \sim x \leftarrow \sim a, \sim b & \sim x \leftarrow \sim a, x \\ \sim x \leftarrow b, \sim x & \sim x \leftarrow x, \sim x \end{array}$$


A rule $x \leftarrow A$ is minimal, if there is no rule $x \leftarrow B$ with $B \subset A$. For determining the supported value, you only need

to take minimal rules into account. The justification frame $U(\mathcal{JF})$ has exactly the following minimal rules:

$$\begin{array}{lll} x \leftarrow a, x & x \leftarrow b, \sim x & x \leftarrow a, b \\ \sim x \leftarrow \sim a, \sim b & \sim x \leftarrow \sim a, x & \sim x \leftarrow \sim b, \sim x \end{array}$$

Of course, it contains many non-minimal rules, for example $x \leftarrow a, b, x$.

The justifications in the original system containing x as an internal nodes are exactly the following:



Assume from now on we are working under \mathcal{B}_{st} . The value of the upper left justification for x is \mathbf{f} in every interpretation. The values of the other justifications for x are at most \mathbf{u} in \mathcal{B}_{st} -models. If it would be \mathbf{t} , then the value of these justifications for x is equal to the value of $\sim x$, which is \mathbf{f} .

By taking the ultimate justification frame, the minimal rule $x \leftarrow a, b$ is added and the minimal rule $\sim x \leftarrow x, \sim x$ is removed. This allows for the justification $a \leftarrow x \rightarrow b$. If the interpretation of a and b is \mathbf{t} , then the value of this justification for x is \mathbf{t} . Therefore, $(\{a, b, x\}, \{a, b, x\})$ is an ultimate stable model, while not a stable model. Note that the lower right justification is not a justification in $U(\mathcal{JF})$. If it would be, then this is a true justification for $\sim x$ contradicting the consistency.

It can be seen that the construction adds rules to \mathcal{JF} in two cases. For the first type, if $x \leftarrow B$ is a rule in R with $B \subseteq A$, then if B is sufficient to derive x , clearly so is A . The second type of rule addition essentially performs some sort of case splitting. It states that if a set of facts A can be used with either y or $\sim y$ to derive x , then the essence for deriving x is the set A itself. In that case, the rule $x \leftarrow A$ is added to the ultimate frame. It turns out that this rule of case splitting is indeed sufficient to reconstruct the ultimate semantics in JT. This is formalised in the main theorem of this section:

Theorem 4. For any frame \mathcal{JF} , $A_{U(\mathcal{JF})} = U(O_{\mathcal{JF}})$.

An immediate corollary is, for instance that the set of stable models of $U(\mathcal{JF})$ equals the set of ultimate stable fixpoints of $O_{\mathcal{JF}}$, and similarly for other semantics. Recall that, in the context of lattices with the subset order, which is what we are concerned with here, the ultimate approximator is defined as follows [Denecker *et al.*, 2004]:

$$U(O)(I_1, I_2) = \left(\bigcap_{I_1 \subseteq K \subseteq I_2} O(K), \bigcup_{I_1 \subseteq K \subseteq I_2} O(K) \right). \quad (1)$$

The proof of Theorem 4 makes use of the following intermediate results.

Lemma 5. Let \mathcal{I} be an interpretation and $x \in \mathcal{F}_d$.

If $O_{\mathcal{JF}}(\mathcal{I}')(x) = \mathbf{t}$ (respectively \mathbf{f}) for all exact interpretations \mathcal{I}' with $\mathcal{I}' \geq_p \mathcal{I}$, then $A_{U(\mathcal{JF})}(\mathcal{I})(x) = \mathbf{t}$ (resp. \mathbf{f}).

Proof sketch. Take $X = \{y \in \mathcal{F}_d \mid \mathcal{I}(y) = \mathbf{u}\}$ and let $\mathcal{I} = (I_1, I_2)$. We prove for all $Y \subseteq X$ and all complete consistent sets A over $X \setminus Y$ that $x \leftarrow \{\mathbf{t}\} \cup I_1 \cup \sim(\mathcal{F}_+ \setminus I_2) \cup A$ is a rule in $U(\mathcal{JF})$. If $Y = X$, then we get that $x \leftarrow \{\mathbf{t}\} \cup I_1 \cup \sim(\mathcal{F}_+ \setminus I_2)$ is a rule in $U(\mathcal{JF})$ such that its body is true in \mathcal{I} , which completes the proof. Our claim is proved by transfinite induction on the size of Y . \square

Combining this lemma with Eq. (1) of the ultimate approximator immediately yields that the operator $A_{U(\mathcal{JF})}(\mathcal{I})$ is at least as precise as the ultimate approximator of $O_{\mathcal{JF}}$.

Lemma 6. For all \mathcal{I} we have $U(O_{\mathcal{JF}})(\mathcal{I}) \leq_p A_{U(\mathcal{JF})}(\mathcal{I})$.

Since the ultimate approximator is the most precise approximator of any given operator, all that is left to prove, to indeed obtain Theorem 4 is that $A_{U(\mathcal{JF})}$ indeed approximates $O_{\mathcal{JF}}$. That is the content of the last lemma.

Lemma 7. $A_{U(\mathcal{JF})}$ is an approximator of $O_{\mathcal{JF}}$.

6 Conclusion

In this paper, we presented a general mechanism to translate justification frames into approximating operators and showed that this transformation preserves all semantics the two formalisms have in common. The correspondence we established provides ample opportunity for future work and in fact probably generates more questions than it answers.

By embedding JT in AFT, JT gets access to a rich body of theoretical results developed for AFT, but of course said results are only directly applicable to branch evaluations that have a counterpart in AFT. A question that immediately arises is whether results such as stratification results also apply to other branch evaluations, and which assumptions on branch evaluations would be required for that. Another question that pops up on the JT side is whether concepts such as *groundedness* [Bogaerts *et al.*, 2015] can be transferred.

On the AFT side, this embedding calls for a general algebraic study of *explanations*. Indeed, for certain approximators, namely those that “come from” a justification frame, our results give us a method for answering certain *why* questions in a graph-based manner (with justifications). Lifting this notion of explanation to general approximators would benefit domains of logics that are covered by AFT but not by JT, such as auto-epistemic logic [Moore, 1985] and default logic [Reiter, 1980].

A last question that emerges naturally is how nesting of justification frames, as defined by Denecker *et al.* [2015] fits into this story, and whether it can give rise to notions of nested operators on the AFT side.

Acknowledgements

Many thanks to Maurice Bruynooghe for the valuable comments and suggestions on draft versions. This research received partial funding from the FWO Flanders project G0B2221N.

References

- [Bi *et al.*, 2014] Yi Bi, Jia-Huai You, and Zhiyong Feng. A generalization of approximation fixpoint theory and application. In *Proceedings of RR*, pages 45–59, 2014.
- [Bogaerts and Cruz-Filipe, 2018] Bart Bogaerts and Luís Cruz-Filipe. Fixpoint semantics for active integrity constraints. *AIJ*, 255:43–70, 2018.
- [Bogaerts and Cruz-Filipe, 2021] Bart Bogaerts and Luís Cruz-Filipe. Stratification in approximation fixpoint theory and its application to active integrity constraints. *ACM Trans. Comput. Logic*, 22(1), January 2021.
- [Bogaerts and Van den Broeck, 2015] Bart Bogaerts and Guy Van den Broeck. Knowledge compilation of logic programs using approximation fixpoint theory. *TPLP*, 15(4–5):464–480, 2015.
- [Bogaerts and Weinzierl, 2018] Bart Bogaerts and Antonius Weinzierl. Exploiting justifications for lazy grounding of answer set programs. In *Proceedings of IJCAI*, pages 1737–1745, 2018.
- [Bogaerts *et al.*, 2015] Bart Bogaerts, Joost Vennekens, and Marc Denecker. Grounded fixpoints and their applications in knowledge representation. *AIJ*, 224:51–71, 2015.
- [Cabalar *et al.*, 2014] Pedro Cabalar, Jorge Fandinno, and Michael Fink. Causal graph justifications of logic programs. *TPLP*, 14(4–5):603–618, 2014.
- [Charalambidis *et al.*, 2018] Angelos Charalambidis, Panos Rondogiannis, and Ioanna Symeonidou. Approximation fixpoint theory and the well-founded semantics of higher-order logic programs. *TPLP*, 18(3-4):421–437, 2018.
- [Damásio *et al.*, 2013] Carlos Viegas Damásio, Anastasia Analyti, and Grigoris Antoniou. Justifications for logic programming. In *Proceedings of LPNMR*, pages 530–542, 2013.
- [Denecker and De Schreye, 1993] Marc Denecker and Danny De Schreye. Justification semantics: A unifying framework for the semantics of logic programs. In *Proceedings of LPNMR*, pages 365–379, 1993.
- [Denecker *et al.*, 2000] Marc Denecker, Victor Marek, and Mirosław Truszczyński. Approximations, stable operators, well-founded fixpoints and applications in nonmonotonic reasoning. In *Logic-Based Artificial Intelligence*, pages 127–144, 2000.
- [Denecker *et al.*, 2003] Marc Denecker, Victor Marek, and Mirosław Truszczyński. Uniform semantic treatment of default and autoepistemic logics. *AIJ*, 143(1):79–122, 2003.
- [Denecker *et al.*, 2004] Marc Denecker, Victor Marek, and Mirosław Truszczyński. Ultimate approximation and its application in nonmonotonic knowledge representation systems. *Information and Computation*, 192(1):84–121, 2004.
- [Denecker *et al.*, 2015] Marc Denecker, Gerhard Brewka, and Hannes Strass. A formal theory of justifications. In *Proceedings of LPNMR*, pages 250–264, 2015.
- [Fages, 1990] F. Fages. A New Fixpoint Semantics for General Logic Programs Compared with the Well-Founded and the Stable Model Semantics. In *ICLP*, page 443. MIT Press, 1990.
- [Fitting, 2002] Melvin Fitting. Fixpoint semantics for logic programming — A survey. *Theoretical Computer Science*, 278(1-2):25–51, 2002.
- [Gebser *et al.*, 2009] Martin Gebser, Roland Kaminski, Benjamin Kaufmann, and Torsten Schaub. On the implementation of weight constraint rules in conflict-driven ASP solvers. In *ICLP*, pages 250–264, 2009.
- [Lapauw *et al.*, 2020] Ruben Lapauw, Maurice Bruynooghe, and Marc Denecker. Improving parity game solvers with justifications. In *Proceedings of VMCAI*, pages 449–470, 2020.
- [Marynissen *et al.*, 2018] Simon Marynissen, Niko Passchyn, Bart Bogaerts, and Marc Denecker. Consistency in justification theory. In *Proceedings of NMR*, pages 41–52, 2018.
- [Marynissen *et al.*, 2020] Simon Marynissen, Bart Bogaerts, and Marc Denecker. Exploiting game theory for analysing justifications. *Theory Pract. Log. Program.*, 20(6):880–894, 2020.
- [Moore, 1985] Robert C. Moore. Semantical considerations on nonmonotonic logic. *AIJ*, 25(1):75–94, 1985.
- [Reiter, 1980] Raymond Reiter. A logic for default reasoning. *AIJ*, 13(1-2):81–132, 1980.
- [Schulz and Toni, 2013] Claudia Schulz and Francesca Toni. ABA-based answer set justification. *TPLP*, 13(4–5-Online-Supplement), 2013.
- [Strass, 2013] Hannes Strass. Approximating operators and semantics for abstract dialectical frameworks. *AIJ*, 205:39–70, 2013.
- [Tarski, 1955] Alfred Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 1955.
- [Vennekens *et al.*, 2006] Joost Vennekens, David Gilis, and Marc Denecker. Splitting an operator: Algebraic modularity results for logics with fixpoint semantics. *ACM Trans. Comput. Log.*, 7(4):765–797, 2006.
- [Vennekens *et al.*, 2007] Joost Vennekens, Maarten Mariën, Johan Wittocx, and Marc Denecker. Predicate introduction for logics with a fixpoint semantics. Parts I and II. *Fundamenta Informaticae*, 79(1-2):187–227, 2007.