Proceedings, Fifteenth International Conference on
Principles of Knowledge Representation and Reasoning (KR 2016)

Characterizing Equivalence Notions for Labelling-Based Semantics

Ringo Baumann
Computer Science Institute
Leipzig University, Germany

Abstract

A central question in knowledge representation is the following: given some knowledge representation formalism, is it possible, and if so how, to simplify parts of a knowledge base without affecting its meaning, even in the light of additional information? The term *strong equivalence* was coined in the literature, i.e. strongly equivalent knowledge bases can be locally replaced by each other in a bigger theory without changing the semantics of the latter. In contrast to classical (monotone) logics where standard and strong equivalence coincide, it is possible to find ordinary but not strongly equivalent objects for any nonmonotonic formalism available in the literature. This paper addresses these questions in the context of abstract argumentation theory. Much effort has been spent to characterize several argumentation tailored equivalence notions w.r.t. extension-based semantics. In recent times labelling-based semantics have received increasing attention, for example in connection with algorithms computing extensions, proof procedures, dialogue games, dynamics in argumentation as well as belief revision in general. Of course, equivalence notions allowing for replacements are of high interest for the mentioned topics. In this paper we provide kernel-based characterization theorems for semantics based on complete labellings as well as admissible labellings w.r.t. eight different equivalence notions including the aforementioned most prominent one, namely strong equivalence.

Introduction

In the last 20 years *formal argumentation* has become a popular research area in AI. One main reason for this is its numerous fields of application ranging from nonmonotonic reasoning, multi-agent systems to analysis tool for debates or dialogues in general (see (Bench-Capon and Dunne 2007; Rahwan and Simari 2009) for excellent overviews). Dung's well-studied *abstract argumentation frameworks* (AFs) play a dominant role in this field (Dung 1995). Here, arguments as well as attacks between them are treated as undefined primitives, i.e. the internal structure of arguments is not considered. This allows one to represent AFs as directed graphs. The major focus in abstract argumentation is on resolving conflicts, or more precisely, on the question of how to determine acceptable sets of arguments. To this end a variety of *semantics* have been introduced, each of them specifying different criteria for being acceptable. There are two main approaches to argumentation semantics, namely *extension-based* and *labelling-based* semantics (cf. (Baroni, Caminada, and Giacomin 2011) for an in-depth overview). In contrast to the latter, extension-based semantics have been studied in detail including central questions like *intertranslatability* (Dvorák and Woltran 2014), *realizability* (Dunne et al. 2015) as well as *replaceability* (Oikarinen and Woltran 2011), among others. In this paper we consider replaceability for labelling-based semantics. More precisely, we studied the following research question:

> Is it possible, and if so, under which conditions, to locally replace parts of a given AF, s.t. the modified version and the initial framework cannot be semantically distinguished, even in the light of additional information?

For this task w.r.t. other formalisms the term *strong equivalence* was coined in the literature, i.e. strongly equivalent theories can be replaced by each other within a bigger theory without changing the meaning of the latter. In contrast to classical (monotone) logics where standard and strong equivalence coincide, it is possible to find ordinarily but not strongly equivalent theories for any nonmonotonic formalism available in the literature. Consequently, much effort has been devoted to characterizing strong equivalence for nonmonotonic formalisms, such as logic programs (Lifschitz, Pearce, and Valverde 2001), causal theories (Turner 2004), default logic (Turner 2001) and nonmonotonic logics in general (Truszczynski 2006; Baumann and Strass 2016).

The characterization theorems in case of abstract argumentation (Oikarinen and Woltran 2011) are quite different from those for the aforementioned formalisms since being strongly equivalent can be decided syntactically in abstract argumentation. More precisely, the authors introduced the notion of a *kernel* of an AF F, which is (informally speaking) a subgraph of F where certain attacks are deleted, and showed that syntactical identity of suitably chosen kernels characterizes strong equivalence w.r.t. the considered semantics. Later it was pointed out that in many argumentation scenarios the potentially occurring type of modification can be anticipated and, more importantly, does not range over *arbitrary expansions* as required for strong equivalence (Baumann 2012; 2014a). This applies, for instance, if we use argumentation theory for the purpose of nonmonotonic entailment, so-called *instantiation-based argumentation* (Caminada and Amgoud 2007), where AFs are built from an underlying knowledge base. For instance, so-called *normal* and *local expansions (deletions)* correspond to re-instantiations if a new piece of information is added (deleted) or if we change to a less (more) restrictive notion of attack. It turned out that many of the corresponding equivalence notions can be decided syntactically, as well (Baumann 2012; 2014a).

As already mentioned (and to the best of our knowledge) all existing characterization theorems are stated in terms of extension-based semantics. Studying equivalence notions w.r.t. labelling-based semantics has both theoretical as well as practical motivations. Among others we can mention,

- *Extension-based vs. Labelling-based Semantics*: It is already known that many semantics establish a one-to-one correspondence between their extension-based and labelling-based versions. This means, any labelling is associated with exactly one extension and vice versa. It is not immediately apparent whether this property guarantees that there is a coincidence of the extension-based and labelling-based equivalence notions. In this paper we give a negative answer which makes characterizing labelling-based equivalence notions a non-trivial task.
- *Labelling-based Algorithms*: As a matter of course, characterization theorems can be used to refine existing algorithms via applying a preprocessing simplification step. Furthermore, the obtained insights may even give rise to new computational procedures.
- *Implicit vs. Explicit Information*: Analogously to other nonmonotonic formalisms we have that possessing the same labellings (explicit information) is not sufficient for semantical indistinguishability w.r.t. further evaluations. It is an interesting theoretical task to precisely determine the gap between implicit and explicit information.

The structure and the main contributions of this paper are as follows. We start by reviewing the necessary background in abstract argumentation frameworks, including extension-based and labelling-based semantics, the whole landscape of notions of equivalence and kernels as well as already existing characterization results for extension-based semantics. In particular, we consider 8 different equivalence notions w.r.t. 8 prominent labelling-based semantics. In total, we provide 61 kernel-based characterization results distributed over three sections. In effect, similarly to extension-based semantics, almost all labelling-based equivalence notions can be decided syntactically. We start with some preliminary results showing that update, deletion as well as local deletion equivalence collapse to identity for any considered labelling-based semantics. Then we turn to the two main as well as considerably more difficult parts of the paper.

In the first one it is shown that many of the already existing characterization results can be transferred to labelling-based semantics. For instance, for all complete-based semantics we have that the labelling-based as well as extension-based versions of strong equivalence coincide. The second main part provides examples showing the non-coincidence of the remaining equivalence notions. Furthermore, we present kernel-based characterizations. For example, we show that strong equivalence w.r.t. admissible labellings is characterized by the complete (instead of the admissible) kernel. Furthermore, the more sophisticated σ-*-kernels are not needed to characterize strong expansion as well as normal deletion equivalence in their labelling-based versions. In the penultimate section we summarize the achieved results and compare them with their extension-based versions. Finally, we discuss related work and present our conclusions. Due to the limited space we omit many proofs. Nevertheless, in almost all cases we included some short comments indicating how to prove the statement in question.

Formal Preliminaries

We start with the necessary background on abstract argumentation. An *argumentation framework* (AF) is a pair $F = (A, R)$ where A, the set of arguments, is a finite subset of a fixed infinite background set $\mathcal{U}$, and $R \subseteq A \times A$. The set of all finite AFs is denoted by $\mathscr{A}_{\text{fin}}$. We say *a attacks b*, or b is *defeated* by a in F whenever $(a, b) \in R$. Furthermore, an argument $a \in A$ is *defended* by a set $A' \subseteq A$ in F if for each $b \in A$ with $(b, a) \in R$, b is attacked by some $a' \in A'$ in F. For a set $E \subseteq A$ we use $R_F^+(E)$ or simply, E^+ for $\{b \mid (a, b) \in R, a \in E\}$. If $G = (B, S)$, we use $A(G)$ as well as $R(G)$ to refer to the first or second component of G, i.e. B or S, respectively. The set $L(G) = \{a \in A(G) \mid (a, a) \in R(G)\}$ contains all self-defeating arguments. A set E is said to be *conflict-free* w.r.t. $F = (A, R)$ if there are no $a, b \in A$, s.t. $(a, b) \in R$. We denote the set of all conflict-free sets by $cf(F)$. In the following we introduce the two main versions of argumentation semantics, namely the extension-based approach and the labelling-based approach (cf. (Baroni, Caminada, and Giacomin 2011) for an excellent overview).

Extension-based Semantics

An *extension-based semantics* $\mathcal{E}_\sigma : \mathscr{A}_{\text{fin}} \to 2^{2^{\mathcal{U}}}$ is a function which assigns to any AF $F = (A, R)$ a set of sets of arguments denoted by $\mathcal{E}_\sigma(F) \subseteq 2^A$. Each one of them, a so-called *σ-extension* E, is considered to be acceptable with respect to F. We consider eight prominent semantics, namely admissible, complete, preferred, semi-stable, stable, stage, grounded, ideal and eager semantics (abbreviated by ad, co, pr, ss, stb, gr, il and eg respectively). From now on we use Φ for the set of all eight semantics mentioned in this paragraph.

Definition 1. Let $F = (A, R)$ be an AF and $E \in 2^A$.

1. $E \in \mathcal{E}_{ad}(F)$ iff $E \in cf(F)$ and E defends all its elements,
2. $E \in \mathcal{E}_{co}(F)$ iff $E \in \mathcal{E}_{ad}(F)$ and for any $a \in A$ defended by E in F, $a \in E$,
3. $E \in \mathcal{E}_{pr}(F)$ iff $E \in \mathcal{E}_{co}(F)$ and there is no $E' \in \mathcal{E}_{ad}(F)$, s.t. $E \subset E'$,
4. $E \in \mathcal{E}_{ss}(F)$ iff $E \in \mathcal{E}_{co}(F)$ and there is no $E' \in \mathcal{E}_{co}(F)$, s.t. $E \cup E^+ \subset E' \cup E'^+$,
5. $E \in \mathcal{E}_{stb}(F)$ iff $E \in \mathcal{E}_{co}(F)$ and $E \cup E^+ = A$,

6. $E \in \mathcal{E}_{gr}(F)$ iff $E \in \mathcal{E}_{co}(F)$ and there is no $E' \in \mathcal{E}_{co}(F)$, s.t. $E' \subset E$,
7. $E \in \mathcal{E}_{il}(F)$ iff $E \in \mathcal{E}_{co}(F)$, $E \subseteq \bigcap \mathcal{E}_{pr}(F)$ and there is no $E' \in \mathcal{E}_{co}(F)$ satisfying $E' \subseteq \bigcap \mathcal{E}_{pr}(F)$ s.t. $E \subset E'$,
8. $E \in \mathcal{E}_{eg}(F)$ iff $E \in \mathcal{E}_{co}(F)$, $E \subseteq \bigcap \mathcal{E}_{ss}(F)$ and there is no $E' \in \mathcal{E}_{co}(F)$ satisfying $E' \subseteq \bigcap \mathcal{E}_{ss}(F)$ s.t. $E \subset E'$. $\diamondsuit$

The following AFs exemplify the different extension-based semantics introduce above and serve as running examples in this paper.

Example 1. Consider the AFs F, G and H.

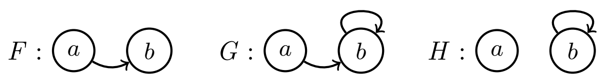

For all $\sigma \in \Phi, \tau \in \Phi \setminus \{stb\}$ we find $\mathcal{E}_\sigma(F) = \mathcal{E}_\sigma(G) = \mathcal{E}_\tau(H)$. Furthermore, $\mathcal{E}_{ad}(F) = \{\emptyset, \{a\}\}$ and for any $\sigma \in \Phi \setminus \{ad\}$, $\mathcal{E}_\sigma(F) = \{\{a\}\}$. Finally, and importantly, $\mathcal{E}_{stb}(H) = \emptyset$. $\diamondsuit$

Labelling-based Semantics

A *labelling-based semantics* $\mathcal{L}_\sigma : \mathscr{A}_{\text{fin}} \to 2^{\left(2^{\mathcal{U}}\right)^3}$ is a function which assigns to any AF $F = (A, R)$ a set of triples of sets of arguments denoted by $\mathcal{L}_\sigma(F) \subseteq \left(2^A\right)^3$. Each one of them, a so-called *σ-labelling* $L = (I, O, U)$ indicates that arguments in I, O or U are considered to be *accepted (in)*, *rejected (out)* or *undecided* with respect to F. We use L^{I} (or $L^{\mathrm{I}}(a)$) to refer to (a is an element of) the first component of the labelling L. Analogously for L^{O} and L^{U}. Furthermore, for two labellings L_1, L_2 we write $L_1 \sqsubseteq L_2$ iff $L_1^{\mathrm{I}} \subseteq L_2^{\mathrm{I}}$ and $L_1^{\mathrm{O}} \subseteq L_2^{\mathrm{O}}$, and as usual $L_1 \sqsubset L_2$ iff $L_1 \sqsubseteq L_2$ and $L_1 \neq L_2$. Finally, given a set of labellings $\mathcal{S}$ we define

$$\sqcap \mathcal{S} = \left(\bigcap_{L \in \mathcal{S}} L^{\mathrm{I}}, \bigcap_{L \in \mathcal{S}} L^{\mathrm{O}}, A \setminus \left(\bigcap_{L \in \mathcal{S}} L^{\mathrm{I}} \cup \bigcap_{L \in \mathcal{S}} L^{\mathrm{O}} \right) \right).$$

We proceed with the central notions of *admissible* as well as *complete labellings*.

Definition 2. A labelling L of an AF $F = (A, R)$ is an *admissible labelling* iff for any $a \in A$,

1. $L^{\mathrm{I}}(a) \to (\forall b : (b, a) \in R \to L^{\mathrm{O}}(b))$ and
2. $L^{\mathrm{O}}(a) \to (\exists b : (b, a) \in R \wedge L^{\mathrm{I}}(b))$.

If it additionally satisfies

3. $L^{\mathrm{I}}(a) \leftarrow (\forall b : (b, a) \in R \to L^{\mathrm{O}}(b))$ and
4. $L^{\mathrm{O}}(a) \leftarrow (\exists b : (b, a) \in R \wedge L^{\mathrm{I}}(b))$

we call it a *complete labelling*. $\diamondsuit$

We introduce the labelling-based counterparts to the extension-based semantics presented in Definition 1.

Definition 3. Let $F = (A, R)$ be an AF and $L \in \left(2^A\right)^3$.

1. $L \in \mathcal{L}_{ad}(F)$ iff L is a admissible labelling of F,
2. $L \in \mathcal{L}_{co}(F)$ iff L is a complete labelling of F,
3. $L \in \mathcal{L}_{pr}(F)$ iff $L \in \mathcal{L}_{co}(F)$ and there is no $M \in \mathcal{L}_{co}(F)$, s.t. $L^{\mathrm{I}} \subset M^{\mathrm{I}}$,
4. $L \in \mathcal{L}_{ss}(F)$ iff $L \in \mathcal{L}_{co}(F)$ and there is no $M \in \mathcal{L}_{co}(F)$, s.t. $M^{\mathrm{U}} \subset L^{\mathrm{U}}$,
5. $L \in \mathcal{L}_{stb}(F)$ iff $L \in \mathcal{L}_{co}(F)$ and $L^{\mathrm{U}} = \emptyset$,
6. $L \in \mathcal{L}_{gr}(F)$ iff $L \in \mathcal{L}_{co}(F)$ and there is no $M \in \mathcal{L}_{co}(F)$, s.t. $M^{\mathrm{I}} \subset L^{\mathrm{I}}$,
7. $L \in \mathcal{L}_{il}(F)$ iff $L \in \mathcal{L}_{co}(F)$, $L \sqsubseteq \sqcap \mathcal{L}_{pr}(F)$ and there is no $M \in \mathcal{L}_{co}(F)$, s.t. $M \sqsubseteq \sqcap \mathcal{L}_{pr}(F)$ and $L \sqsubset M$,
8. $L \in \mathcal{L}_{eg}(F)$ iff $L \in \mathcal{L}_{co}(F)$, $L \sqsubseteq \sqcap \mathcal{L}_{ss}(F)$ and there is no $M \in \mathcal{L}_{co}(F)$, s.t. $M \sqsubseteq \sqcap \mathcal{L}_{ss}(F)$ and $L \sqsubset M$. $\diamondsuit$

Example 2 (Example 1 cont.). In case of labelling-based semantics we observe certain differences. More precisely, $\mathcal{L}_\sigma(F) = \mathcal{L}_\sigma(G)$ but $\mathcal{L}_\sigma(F) \neq \mathcal{L}_\sigma(H)$ for any $\sigma \in \Phi$. Furthermore, $\mathcal{L}_{ad}(F) = \{(\emptyset, \emptyset, \{a, b\}), (\{a\}, \emptyset, \{b\}), (\{a\}, \{b\}, \emptyset)\}$ and $\mathcal{L}_{ad}(H) = \{(\emptyset, \emptyset, \{a, b\}), (\{a\}, \emptyset, \{b\})\}$. For the remaining $\sigma \in \Phi \setminus \{ad\}$ we obtain $\mathcal{L}_\sigma(F) = \{(\{a\}, \{b\}, \emptyset)\}$ and for $\tau \in \Phi \setminus \{ad, stb\}$, $\mathcal{L}_\tau(H) = \{(\{a\}, \emptyset, \{b\})\}$. Finally, $\mathcal{L}_{stb}(H) = \emptyset$. $\diamondsuit$

Basic Properties and Relations

In the following we list well known properties and relations between semantics which are frequently used in the proofs. For more details and explanations confer (Dung 1995; Caminada and Gabbay 2009; Caminada 2011; Caminada and Pigozzi 2011; Baumann and Spanring 2015).

Fact 1. *Let $X \in \{\mathcal{L}, \mathcal{E}\}$. For any AF F we have*

1. $X_\sigma(F) \subseteq X_{co}(F)$ *for all* $\sigma \in \Phi \setminus \{ad\}$,
2. $|X_\sigma(F)| \geq 1$ *for all* $\sigma \in \Phi \setminus \{stb\}$ *and*
3. $|X_{gr}(F)| = |X_{il}(F)| = |X_{eg}(F)| = 1$.

Fact 2. *Given an AF $F = (A, R)$ and $E \subseteq A$. We write $E^{\mathcal{L}}$ for $(E, E^+, A \setminus (E \cup E^+))$. For all $\sigma \in \Phi$ we have,*

1. *If $L \in \mathcal{L}_\sigma(F)$, then $L^I \in \mathcal{E}_\sigma(F)$,*
2. *If $E \in \mathcal{E}_\sigma(F)$, then $E^{\mathcal{L}} \in \mathcal{L}_\sigma(F)$ and*
3. *Obviously, $(E^{\mathcal{L}})^I = E$.*

We point out that the properties mentioned in Fact 2 do not ensure that there is a one-to-one correspondence between σ-labellings and σ-extensions. This desirable feature (which would indeed justify the terms σ-labellings and σ-extensions) is given if additionally, labellings are uniquely determined by their in-labelled arguments.

Fact 3. *Given an AF $F = (A, R)$ and a set $E \subseteq A$. Let $\sigma \in \Phi \setminus \{ad\}$. We have,*

1. *For any $L, M \in \mathcal{L}_\sigma(F), L^I = M^I$ iff $L = M$,*
2. *Given $L \in \mathcal{L}_\sigma(F)$, then $(L^I)^{\mathcal{L}} = L$ and*
3. $|\mathcal{L}_\sigma(F)| = |\mathcal{E}_\sigma(F)|$.

Please note that admissible labellings are excluded from Fact 3. The AF F handled in Examples 1 and 2 shows that this is no coincidence.

Replaceability and Characterization Theorems

In this section we review typical dynamic scenarios and introduce their corresponding equivalence notions. Furthermore, we present already existing characterization results for extension-based semantics.

Dynamic Scenarios and Equivalence Notions

There are two main classes of dynamic scenarios, namely *expansions* and *deletions* (see Figure 1). Both of them can be further divided in *normal* and *local* versions. These scenarios are motivated by real-world argumentation as well as instantiation-based argumentation (cf. (Caminada and Amgoud 2007; Baumann and Brewka 2010; Baumann 2014a) as well as (Baumann 2014b, Section 2.1.4) for more explanations).

Definition 4. Given an AF $F = (A, R)$, a set of arguments B and a set of attacks S as well as a further AF H. The AF

$$G = (F \setminus [B, S]) \cup H = \big((A, R \setminus S)|_{A \setminus B}\big) \cup H$$

is called an *update* of F (for short, $F \asymp_U G$). An update is called a

1. *deletion* ($F \succeq_D G$) iff $H = \emptyset$,
2. *normal deletion* ($F \succeq_{ND} G$) iff $F \succeq_D G$ and $S = \emptyset$,
3. *local deletion* ($F \succeq_{LD} G$) iff $F \succeq_D G$ and $B = \emptyset$,
4. *expansion* ($F \preceq_E G$) iff $B = S = \emptyset$,
5. *normal expansion* ($F \preceq_N G$) iff $F \preceq_E G = (C, T)$ and $\forall ab\ ((a, b) \in T \setminus R \rightarrow a \in C \setminus A \vee b \in C \setminus A)$,
6. *strong expansion* ($F \preceq_S G$) iff $F \preceq_N G = (C, T)$ and $\forall ab\ ((a, b) \in T \setminus R \rightarrow \neg(a \in C \setminus A \wedge b \in A))$,
7. *local expansion* ($F \preceq_L G$) iff $F \preceq_E G = (C, T)$ and $A = C$. $\Diamond$

Consider again Example 1. The AF G is a local expansion of F and H is a local deletion of G.

Notions of Equivalence

We now introduce the corresponding equivalence notions (cf. (Baumann and Strass 2015, Section 3.8) for chronological order). If two (possibly syntactically different) AFs are equivalent, then they share the same implicit information w.r.t. further expansions or deletions, respectively. This means, they cannot be semantically distinguished in any suitable future scenario and thus, can be replaced by each other without loss of (semantical) information. The first paper in this line of work was (Oikarinen and Woltran 2011) engaged with characterizing *strong equivalence*. For the sake of clarity and comprehensibility we use the term *expansion equivalence* to indicate that arbitrary expansions are allowed.

Definition 5. Given a semantics σ and let $X \in \{\mathcal{L}, \mathcal{E}\}$. Two AFs F and G are

1. *standard equivalent* w.r.t. X_σ ($F \equiv^{X_\sigma} G$) iff $X_\sigma(F) = X_\sigma(G)$,
2. *update equivalent* w.r.t. X_σ ($F \equiv_U^{X_\sigma} G$) iff for any pair $[B, S]$ and any AF H we have: $(F \setminus [B, S]) \cup H \equiv^{X_\sigma} (G \setminus [B, S]) \cup H$,
3. *deletion equivalent* w.r.t. X_σ ($F \equiv_D^{X_\sigma} G$) iff for any pair $[B, S]$ we have: $F \setminus [B, S] \equiv^{X_\sigma} G \setminus [B, S]$,
4. *normal deletion equivalent* w.r.t. X_σ ($F \equiv_{ND}^{X_\sigma} G$) iff for any set of arguments B we have: $F \setminus [B, \emptyset] \equiv^{X_\sigma} G \setminus [B, \emptyset]$,
5. *local deletion equivalent* w.r.t. X_σ ($F \equiv_{LD}^{X_\sigma} G$) iff for any set of attacks S we have: $F \setminus [\emptyset, S] \equiv^{X_\sigma} G \setminus [\emptyset, S]$,
6. *expansion equivalent* w.r.t. X_σ ($F \equiv_E^{X_\sigma} G$) iff for each AF H we have: $F \cup H \equiv^{X_\sigma} G \cup H$,
7. *normal expansion equivalent* w.r.t. X_σ ($F \equiv_N^{X_\sigma} G$) iff for each AF H, such that $F \preceq_N F \cup H$ and $G \preceq_N G \cup H$ we have: $F \cup H \equiv^{X_\sigma} G \cup H$,
8. *strong expansion equivalent* w.r.t. X_σ ($F \equiv_S^{X_\sigma} G$) iff for each AF H, such that $F \preceq_S F \cup H$ and $G \preceq_S G \cup H$ we have: $F \cup H \equiv^{X_\sigma} G \cup H$,
9. *local expansion equivalent*[1] w.r.t. X_σ ($F \equiv_L^{X_\sigma} G$) iff for each AF H, such that $A(H) \subseteq A(F \cup G)$ we have: $F \cup H \equiv^{X_\sigma} G \cup H$. $\Diamond$

From now on we use $\mathcal{M}$ as a shorthand for $\{L, E, N, S, ND, D, LD, U\}$. Furthermore, we use M-*equivalent* as a placeholder for any equivalence notion defined above. The following figure gives a preliminary overview about interrelations (arising from the definitions) between the introduced equivalence notions. For two equivalence notion Φ and Ψ we have $\Phi \subseteq \Psi$ iff there is a link from Φ to Ψ.

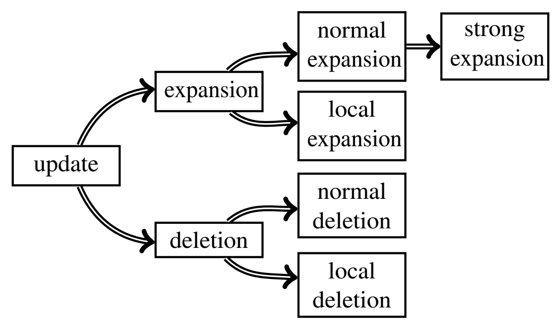

Figure 1: Preliminary Relations

Example 3 (Examples 1, 2 cont.). We mention only a few (non-)relations. We already observed that F and G are standard equivalent w.r.t. all semantics considered in this paper. Furthermore, they can not be distinguished w.r.t. the grounded semantics and local expansion. In contrast, the AF $I = (\{a, b\}, \{(b, a)\})$ shows that they are not local equivalent w.r.t. all considered semantics except the grounded one.

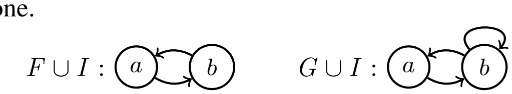

Extension-based Semantics are Characterizable through Kernels

A *kernel* is a function $k : \mathscr{A}_{\text{fin}} \rightarrow \mathscr{A}_{\text{fin}}$ where each $k(F) = F^k$ is obtained from F by deleting certain (redundant) information. A relation $\equiv$ is *characterizable through kernels* if there is a kernel k, s.t. $F \equiv G$ iff $F^k = G^k$. We also say k *characterizes* $\equiv$ or k *is the characterizing kernel of* $\equiv$. It was the main result in (Oikarinen and Woltran 2011) that expansion equivalence is characterizable through kernels.

[1]Note that a suitable AF H is not necessarily a local expansion of F and G in the sense of Definition 4. Nevertheless, we may loosely speak about local expansions.

Later it was shown that almost all extension-based equivalence notions presented in Definition 5 are characterizable through kernels (Baumann 2012; 2014a). In the following we list all relevant kernel definitions.

Definition 6. Given an AF $F = (A, R)$ and a semantics σ. We define σ-kernels $F^{k(\sigma)} = \left(A, R^{k(\sigma)}\right)$ as well as σ-*-kernels $F^{k^*(\sigma)} = \left(A, R^{k^*(\sigma)}\right)$ whereby

$$\begin{aligned}
R^{k(stb)} &= R \setminus \{(a,b) \mid a \neq b, (a,a) \in R\},\\
R^{k(ad)} &= R \setminus \{(a,b) \mid a \neq b, (a,a) \in R,\\
&\quad \{(b,a),(b,b)\} \cap R \neq \emptyset\},\\
R^{k(gr)} &= R \setminus \{(a,b) \mid a \neq b, (b,b) \in R,\\
&\quad \{(a,a),(b,a)\} \cap R \neq \emptyset\},\\
R^{k(co)} &= R \setminus \{(a,b) \mid a \neq b, (a,a),(b,b) \in R\},\\
R^{k^*(ad)} &= R \setminus \{(a,b) \mid a \neq b, ((a,a) \in R \wedge \{(b,a),(b,b)\}\\
&\quad \cap R \neq \emptyset) \vee ((b,b) \in R \wedge \forall c\, ((b,c) \in R \rightarrow\\
&\quad \{(a,c),(c,a),(c,c),(c,b)\} \cap R \neq \emptyset))\},\\
R^{k^*(gr)} &= R \setminus \{(a,b) \mid a \neq b, ((b,b) \in R \wedge \{(a,a),(b,a)\}\\
&\quad \cap R \neq \emptyset) \vee ((b,b) \in R \wedge \forall c\, ((b,c) \in R \rightarrow\\
&\quad \{(a,c),(c,a),(c,c)\} \cap R \neq \emptyset))\},\\
R^{k^*(co)} &= R \setminus \{(a,b) \mid a \neq b, ((a,a),(b,b) \in R)\\
&\quad \vee ((b,b) \in R \wedge (b,a) \notin R \wedge \forall c\, ((b,c) \in R \rightarrow\\
&\quad \{(a,c),(c,a),(c,c),(c,b)\} \cap R \neq \emptyset))\}. \qquad \Diamond
\end{aligned}$$

Example 4 (Example 1 cont.). Let $\sigma \in \{ad, gr, co\}$. We have $G^{k(\sigma)} = G \neq H = H^{k(\sigma)}$. On the other hand, (a, b) has to be deleted in G if we apply σ-*-kernels since all second disjuncts are satisfied. More precisely, $G^{k^*(\sigma)} = H^{k^*(\sigma)} = H$. $\Diamond$

From now on we use $\mathcal{K}$ as a shorthand for the set of all kernels introduced in Definition 6. The following figure gives a comprehensive overview over characterization results (see (Baumann and Woltran 2014; Baumann and Brewka 2015) for recent overviews). The entry k in row M and column σ indicates that $\equiv^{\mathcal{E}_\sigma}_M$ is characterizable through k. The abbreviation id stands for identity map and the question mark represents an open problem. Furthermore, $[1]_m$ ($[2]_m$) means that the characterization theorem m in (Oikarinen and Woltran 2011) ((Baumann 2014a)) is not purely kernel-based.

Example 5 (Examples 1, 4 cont.). Given a semantics $\sigma \in \{ad, pr, il, gr, co\}$. In consideration of Figure 2 we deduce that G and H are semantically distinguishable via arbitrary expansions. In contrast, this endeavour is impossible if we restrict ourselves to strong expansions. Formally, $G \equiv^{\mathcal{E}_\sigma}_S H$ and $G \not\equiv^{\mathcal{E}_\sigma}_E H$. $\Diamond$

Finally, we list useful properties of the introduced kernels.

Fact 4. *Let $k \in \mathcal{K}$ and given two AFs F and G.*

1. $cf(F) = cf(F^k)$, $A(F) = A(F^k)$ *and* $L(F) = L(F^k)$,
2. $\mathcal{E}_\sigma(F) = \mathcal{E}_\sigma\left(F^{k(\sigma)}\right)$ *where* $\sigma \in \{stb, gr, co\}$,
3. $\mathcal{E}_\sigma(F) = \mathcal{E}_\sigma\left(F^{k(ad)}\right)$ *where* $\sigma \in \{ss, eg, ad, pr, il\}$,
4. $F^k = G^k \Rightarrow$ *for any AF* H, $(F \cup H)^k = (G \cup H)^k$,
5. $F^k = G^k \Rightarrow$ *for any set* B, $(F \setminus [B, \emptyset])^k = (G \setminus [B, \emptyset])^k$.

	stb	ss	eg	ad	pr	il	gr	co
L	$[1]_9$	$k(ad)$	$k(ad)$	$k(ad)$	$k(ad)$	$k(ad)$	$[1]_{10}$	$[1]_{11}$
E	$k(stb)$	$k(ad)$	$k(ad)$	$k(ad)$	$k(ad)$	$k(ad)$	$k(gr)$	$k(co)$
N	$k(stb)$	$k(ad)$	$k(ad)$	$k(ad)$	$k(ad)$	$k(ad)$	$k(gr)$	$k(co)$
S	$k(stb)$	$k(ad)$	$k(ad)$	$k^*(ad)$	$k^*(ad)$	$k^*(ad)$	$k^*(gr)$	$k^*(co)$
ND	$[2]_{10}$	?	?	$[2]_{16}$	?	?	$[2]_{16}$	$[2]_{16}$
D	id	id	id	id	id	id	id	id
LD	id	id	id	id	id	id	id	id
U	id	id	id	id	id	id	id	id

Figure 2: Extension-based Characterizations

Preliminary Results

The basic properties listed in Facts 1 and 2 are concerned with single AFs. In this section we compare the set of extensions as well as labellings of different AFs. This will lead to the general insight that being M-equivalent w.r.t. labelling-based semantics implies being M-equivalent w.r.t. extension-based semantics. As a by-product of this analysis we obtain first characterization results for labelling-based semantics. More precisely, analogously to extension-based semantics we have that deletion, local deletion as well as update equivalence w.r.t. $\mathcal{L}_\sigma$ collapse to identity. The following lemma can be shown by combining the assertions listed in Fact 2.

Lemma 5. *Given two AFs F and G. For any $\sigma \in \Phi$,*

$$\mathcal{L}_\sigma(F) = \mathcal{L}_\sigma(G) \Rightarrow \mathcal{E}_\sigma(F) = \mathcal{E}_\sigma(G).$$

Observe that the converse direction of Lemma 5 does not hold in general. The main reason for the invalidity is that AFs may possess the same σ-extensions without sharing the same arguments. This is impossible in case of labelling-based semantics since labellings assign a label to any argument of the considered AF implying that labellings are necessarily different if the frameworks in question possess different arguments. Furthermore, even sharing the same arguments does not ensure the validity of the converse direction. Consider therefore the running examples F and H (Example 1). Although they possess the same complete extensions they yield different complete labellings (cf. Example 2).

Having Lemma 5 at hand we may state (and one may easily prove) the fundamental relation between labelling-based and extension-based equivalence notion.

Theorem 6. *Let $\sigma \in \Phi$ and $M \in \mathcal{M}$. For any AFs F,G,*

$$F \equiv^{\mathcal{L}_\sigma}_M G \Rightarrow F \equiv^{\mathcal{E}_\sigma}_M G.$$

In the rest of the paper we will study in detail in which cases the converse direction holds. We point out, that it is a misconception to expect that the extension-based and labelling-based version of a certain equivalence notion coincide if the considered σ-labellings and σ-extensions are in one-to-one correspondence. This does not hold in general and makes characterizing labelling-based equivalence notions a non-trivial task.

We proceed with the first characterization theorem. Remember that identity is the finest equivalence relation. Furthermore, it is already shown that deletion, local deletion as well as update equivalence w.r.t. $\mathcal{E}_\sigma$ collapse to identity (see Figure 2). Consequently, applying the fundamental relation stated in Theorem 6 we obtain the identical characterization results w.r.t. labelling-based semantics.

Theorem 7. *Let $M \in \{D, LD, U\}$ and $\sigma \in \Phi$. For any two AFs F and G,*

$$F \equiv_M^{\mathcal{L}_\sigma} G \Leftrightarrow F = G.$$

Example 6 (Examples 1,2 cont.). Obviously, $F \neq G$. Hence, $F \not\equiv_{LD}^{\mathcal{L}_\sigma} G$ for any $\sigma \in \Phi$ (Theorem 7). The following local deletions exemplify the predicted non-equivalence. Let $S = \{(a,b)\}$. Consequently, $F \setminus [\emptyset, S] = (\{a,b\}, \emptyset)$ and $G \setminus [\emptyset, S] = (\{a,b\}, \{(a,b)\}) = H$. For any $\sigma \in \Phi$, $(\{a,b\}, \emptyset, \emptyset) \in \mathcal{L}_\sigma(F \setminus [\emptyset, S]) \setminus \mathcal{L}_\sigma(G \setminus [\emptyset, S])$ since b is self-defeating in H.

Coincidence of Extension-based and Labelling-based Equivalence Notions

Two labellings of the same AF are identical if they share the same accepted as well as rejected arguments. If considering labellings of different AFs this assertion remains true as long as they share the same arguments. Due to the definition of complete labellings we may state an even weaker condition for equality.

Proposition 8. *Given two AFs F and G as well as $L \in \mathcal{L}_{co}(F)$ and $M \in \mathcal{L}_{co}(G)$. We have $L = M$ if simultaneously $A(F) = A(G)$, $L^I = M^I$ and $R_F^+(L^I) = R_G^+(M^I)$.*

The assertion does not hold in case of admissible labellings. Consider therefore the following counterexample.

Example 7 (Examples 1, 3 cont.). Consider the AFs G and $G \cup I$ as well as the the labellings $L = (\{a\}, \emptyset, \{b\}) \in \mathcal{L}_{ad}(G)$ and $M = (\{a\}, \{b\}, \emptyset) \in \mathcal{L}_{ad}(G \cup I)$. Obviously, $A(G) = A(G \cup I) = \{a, b\}$, $L^I = M^I = \{a\}$ and $R_G^+(L^I) = R_{G\cup I}^+(M^I) = \{b\}$ but $L \neq M$. $\diamondsuit$

The following proposition states that some kernels possess the property that the range of a conflict-free set remains the same even if we apply them to the considered framework. The proof requires Definition 6 and Fact 4-1. only.

Proposition 9. *Let $k \in \{k(stb), k(ad), k(co)\}$. For any AF F and $E \subseteq cf(F)$ we have, $R_F^+(E) = R_{F^k}^+(E)$.*

The assertion is not valid for the remaining kernels considered in this paper.

Example 8 (Examples 1, 3, 4 cont.). We already observed that for any $\sigma \in \{ad, gr, co\}$, $G^{k^*(\sigma)} = H$. Moreover, although $\{a\} \in cf(G) = cf\left(G^{k^*(\sigma)}\right)$ we observe $R_G^+(\{a\}) = \{b\} \neq \emptyset = R_{G^{k^*(\sigma)}}^+(\{a\})$. Similarly, one may check that the AF $G \cup I$ serves as a counter-example in case of the grounded kernel. $\diamondsuit$

Similarly to Fact 4-2./3. we will show that there are combinations of kernels and semantics σ, s.t. the application of a kernel does not change the set of σ-labellings.

Lemma 10. *For any AF F,*

1. *$\mathcal{L}_\sigma(F) = \mathcal{L}_\sigma\left(F^{k(\sigma)}\right)$ for $\sigma \in \{co, stb, gr\}$ and*
2. *$\mathcal{L}_\tau(F) = \mathcal{L}_\tau\left(F^{k(ad)}\right)$ for $\tau \in \{ss, eg, pr, il\}$.*

Proof. *1. Due to the limited space we exclude a proof for grounded semantics since it is structurally different. Given $\sigma \in \{co, stb\}$. ($\subseteq$) Let $L \in \mathcal{L}_\sigma(F)$. Hence, $L^I \in \mathcal{E}_\sigma(F)$ (Fact 2-1.). Furthermore, $L^I \in \mathcal{E}_\sigma(F^{k(\sigma)})$ (Fact 4-2.). Consequently, $(L^I)^{\mathcal{L}} \in \mathcal{L}_\sigma\left(F^{k(\sigma)}\right)$ (Fact 2-2.). Note that $\left((L^I)^{\mathcal{L}}\right)^I = L^I$ and moreover $\left((L^I)^{\mathcal{L}}\right)^O = R_{F^{k(\sigma)}}^+(L^I) = R_F^+(L^I)$ (Proposition 9). This means, $L = (L^I)^{\mathcal{L}} \in \mathcal{L}_\sigma\left(F^{k(\sigma)}\right)$ (Proposition 8). ($\supseteq$) Given $L \in \mathcal{L}_\sigma(F^{k(\sigma)})$. Thus, $L^I \in \mathcal{E}_\sigma(F^{k(\sigma)})$ (Fact 2-1.). Again, $L^I \in \mathcal{E}_\sigma(F)$ (Fact 4-2.). Consequently, $(L^I)^{\mathcal{L}} \in \mathcal{L}_\sigma(F)$ (Fact 2-2.). Note that $\left((L^I)^{\mathcal{L}}\right)^I = L^I$ and moreover $\left((L^I)^{\mathcal{L}}\right)^O = R_F^+(L^I) = R_{F^{k(\sigma)}}^+(L^I)$ (Proposition 9). Finally, $L = (L^I)^{\mathcal{L}} \in \mathcal{L}_\sigma(F)$ (Proposition 8) concluding the proof.*
2. Similarly. Use Fact 4-3. instead of Fact 4-2. Furthermore, Proposition 8 is applicable since all considered semantics are complete ones.

Now we are prepared for the main theorem of this section. It stipulates that several expansion equivalence relations as well as weaker notions do not distinguish between their labelling-based and extension-based version. This means, kernel-based characterization results (depicted in Figure 1) carry over to labelling-based semantics. In order to save space we formalize the theorem in terms of extension-based/labelling-based equivalence notions. An overview of characterizing kernels is presented in Figure 3, Section Summary and Comparison.

Theorem 11. *Given two AFs F and G. We have,*

1. *$F \equiv_M^{\mathcal{E}_\sigma} G \Leftrightarrow F \equiv_M^{\mathcal{L}_\sigma} G$ for $\sigma \in \Phi \setminus \{ad\}$, $M \in \{E, N\}$,*
2. *$F \equiv_L^{\mathcal{E}_\sigma} G \Leftrightarrow F \equiv_L^{\mathcal{L}_\sigma} G$ for $\sigma \in \{ss, eg, pr, il\}$ and*
3. *$F \equiv_S^{\mathcal{E}_\sigma} G \Leftrightarrow F \equiv_S^{\mathcal{L}_\sigma} G$ for $\sigma \in \{stb, ss, eg\}$.*

Proof. *($\Leftarrow$) Apply Theorem 6.*
($\Rightarrow$) For the proof we use σ and M as a placeholder for a combination mentioned above. Given an AF H, s.t. $F \cup H$ and $G \cup H$ are (normal, strong, local) expansions of F and G as required in Definition 5. Since $F \equiv_M^{\mathcal{E}_\sigma} G$ is assumed we deduce $F^k = G^k$ for the associated kernel k (compare Figure 1). Thus, $(F \cup H)^k = (G \cup H)^k$ (Fact 4-4.). Due to Lemma 10 we obtain $\mathcal{L}_\sigma(F \cup H) = \mathcal{L}_\sigma\left((F \cup H)^k\right) = \mathcal{L}_\sigma\left((G \cup H)^k\right) = \mathcal{L}_\sigma(G \cup H)$. Hence, $F \equiv_M^{\mathcal{L}_\sigma} G$.

Non-Coincidence of Extension-based and Labelling-based Equivalence Notions

We now leave the realm of uniformity of extension-based and labelling-based characterizations. In particular, we show that the admissible kernel (originally introduced to characterize equivalence notions w.r.t. admissible extension-based semantics) does not serve as characterizing kernel for admissible labellings. Furthermore, we will see that the more sophisticated σ-*-kernels are not needed to characterize strong expansion equivalence. Finally, normal deletion equivalence w.r.t. labelling-based semantics can be decided via traditional kernels in contrast to the extension-based versions.

Expansion Equivalence w.r.t. Admissible Labellings

Expansion equivalence as well as its local, normal and strong versions w.r.t. admissible extensions are characterizable through the admissible kernel. The following example shows that this assertion does not hold in case of admissible labellings.

Example 9. Consider the following AFs.

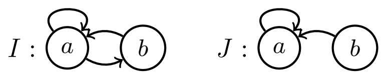

Observe that $I^{k(ad)} = J^{k(ad)} = J$. Consequently, $I \equiv_M^{\mathcal{E}_{ad}} J$ for $M \in \{L, E, N\}$ (cf. Figure 1). On the other hand, $(\{b\}, \emptyset, \{a\}) \in \mathcal{L}_{ad}(J) \setminus \mathcal{L}_{ad}(I)$ because the argument a cannot be undecided in I since it attacks the in-labelled argument b (compare Definition 2). Thus $I \not\equiv_M^{\mathcal{L}_{ad}} J$ for $M \in \{L, E, N, S\}$. $\Diamond$

One of the most surprising results for ourselves is that expansion equivalence as well as its local, normal and strong variant w.r.t. admissible labellings are characterizable through the complete kernel as stated by the following theorem.

Theorem 12. *Given two AFs F and G. We have,*

$$F \equiv_M^{\mathcal{L}_{ad}} G \Leftrightarrow F^{k(co)} = G^{k(co)} \text{ for } M \in \{L, E, N, S\}.$$

The proof of the theorem includes a long case distinction similar to the proof of Theorem 15.

Example 10 (Example 9 cont.). Remember that the complete kernel deletes an attack (a, b) if and only if a and b are self-defeating. Consequently, $I^{k(co)} = I \neq J = J^{k(co)}$ as predicted in Theorem 12. $\Diamond$

Strong Expansion Equivalence for Preferred, Ideal, Grounded and Complete Labellings

In this subsection we will see that σ-*-kernels are not needed to characterize strong expansion equivalence w.r.t. labelling-based preferred, ideal, grounded and complete semantics. Interestingly, we observe a certain symmetry, namely the labelling-based version is characterizable through a σ-kernel if and only if the extension-based version is characterizable through the corresponding σ-*-kernel.

Example 11 (Examples 1, 4 cont.). We already observed that $G^{k^*(\sigma)} = H^{k^*(\sigma)} = H$ for any $\sigma \in \{ad, gr, co\}$. In consideration of Figure 1 we deduce $G \equiv_S^{\mathcal{E}_\sigma} H$ for $\sigma \in \{pr, il, gr, co\}$. On the other hand, $(\{a\}, \{b\}, \emptyset) \in \mathcal{L}_\sigma(G) \setminus \mathcal{L}_\sigma(H)$ because b cannot be out-labelled in H since it is not attacked by an in-labelled argument (compare Definition 2). Therefore, $G \not\equiv_S^{\mathcal{L}_\sigma} H$ for $\sigma \in \{pr, il, gr, co\}$. $\Diamond$

We proceed with a theorem stating that possessing the same σ-kernels is necessary and sufficient for being strong expansion equivalent w.r.t. certain labelling-based semantics.

Theorem 13. *Given two AFs F and G. We have,*

1. $F \equiv_S^{\mathcal{L}_\sigma} G \Leftrightarrow F^{k(ad)} = G^{k(ad)}$ *for* $\sigma \in \{pr, il\}$,
2. $F \equiv_S^{\mathcal{L}_{gr}} G \Leftrightarrow F^{k(gr)} = G^{k(gr)}$ *and*
3. $F \equiv_S^{\mathcal{L}_{co}} G \Leftrightarrow F^{k(co)} = G^{k(co)}$.

Example 12 (Example 9 cont.). We already observed that the AFs I and J possess the same admissible kernel. Consequently, in accordance with Theorem 13 we obtain, $I \equiv_S^{\mathcal{L}_{pr}} J$. We encourage the reader to verify this assertion for some strong expansions. $\Diamond$

Normal Deletion Equivalence

Characterizing normal deletion equivalence in case of extension-based semantics is exceptional in several regards. Remember that normal deletions retract arguments and their corresponding attacks. Firstly, only a few characterization results are achieved, namely in case of stable, admissible, complete and grounded semantics (see (Baumann and Strass 2015, Open Problem 13) for a summary of open problems). Secondly, none of the characterization results is purely kernel-based, i.e. beside the equality of kernels on certain parts of the frameworks further loop- as well as attack-conditions have to be satisfied. Furthermore, quite surprisingly, normal deletion equivalent AFs do not even have to share the same arguments and finally, equivalence classes may have an infinite number of elements.

We will see that characterizing normal deletion equivalence in case of labelling-based semantics is quite different from their extension-based versions. Any labelling-based semantics considered in this paper is characterizable through traditional kernels and thus, do not share any of the features mentioned above. We proceed with a lemma paving the way for the main theorem.

Lemma 14. *Given two AFs F and G. We have*

1. $F^{k(stb)} = G^{k(stb)} \Rightarrow F \equiv_{ND}^{\mathcal{L}_{stb}} G$,
2. $F^{k(ad)} = G^{k(ad)} \Rightarrow F \equiv_{ND}^{\mathcal{L}_\sigma} G$ *for* $\sigma \in \{ss, eg, pr, il\}$,
3. $F^{k(co)} = G^{k(co)} \Rightarrow F \equiv_{ND}^{\mathcal{L}_\sigma} G$ *for* $\sigma \in \{ad, co\}$ *and*
4. $F^{k(gr)} = G^{k(gr)} \Rightarrow F \equiv_{ND}^{\mathcal{L}_{gr}} G$,

Proof. *For the proof we use $\mathcal{L}_\sigma$ and k as a placeholder for a combination mentioned above. Given $F^k = G^k$. Consequently, for any set B, $(F \setminus [B, \emptyset])^k = (G \setminus [B, \emptyset])^k$ (Fact 4-5.) Remember that $\mathcal{L}_\sigma(F \setminus [B, \emptyset]) = \mathcal{L}_\sigma\left((F \setminus [B, \emptyset])^k\right)$ as well as $\mathcal{L}_\sigma(G \setminus [B, \emptyset]) = \mathcal{L}_\sigma\left((G \setminus [B, \emptyset])^k\right)$ (Lemmata 10 and 12). Consequently, we derive $F \equiv_{ND}^{\mathcal{L}_\sigma} G$ since we put no restrictions on B.* $\square$

We now present the main theorem of this section.

Theorem 15. *Given two AFs F and G. We have*

1. $F \equiv^{\mathcal{L}_{stb}}_{ND} G \Leftrightarrow F^{k(stb)} = G^{k(stb)}$,
2. $F \equiv^{\mathcal{L}_{\sigma}}_{ND} G \Leftrightarrow F^{k(ad)} = G^{k(ad)}$ *for* $\sigma \in \{ss, eg, pr, il\}$,
3. $F \equiv^{\mathcal{L}_{\sigma}}_{ND} G \Leftrightarrow F^{k(co)} = G^{k(co)}$ *for* $\sigma \in \{ad, co\}$ *and*
4. $F \equiv^{\mathcal{L}_{gr}}_{ND} G \Leftrightarrow F^{k(gr)} = G^{k(gr)}$,

Proof. ($\Leftarrow$) Apply Lemma 14.
($\Rightarrow$) We show the contrapositive. Assume $A\left(F^{k(\sigma)}\right) \neq A\left(G^{k(\sigma)}\right)$ for $\tau \in \{stb, ad, co, gr\}$. Hence, $A(F) \neq A(G)$ (Fact 4-1.). Since any semantics $\sigma \in \Phi \setminus \{stb\}$ possesses at least one σ-labelling (Fact1-3.) we derive $\mathcal{L}_\sigma(F) \neq \mathcal{L}_\sigma(G)$. Consider now stable labellings. W.l.o.g. let $a \in A(F) \setminus A(G)$. Define $B = (A(F) \cup A(G)) \setminus \{a\}$. We have $G \setminus [B, \emptyset] = (\emptyset, \emptyset) \neq F \setminus [B, \emptyset]$. Obviously, $\{(\emptyset, \emptyset, \emptyset)\} = \mathcal{L}_{stb}(G \setminus [B, \emptyset])$. Since $A(F \setminus [B, \emptyset]) = \{a\}$ we derive $F \not\equiv^{\mathcal{L}_\sigma}_{ND} G$ for any $\sigma \in \Phi$. From now on we assume $A(F) = A(G)$. Suppose now $L\left(F^{k(\tau)}\right) \neq L\left(G^{k(\tau)}\right)$ for $\tau \in \{stb, ad, co, gr\}$. Without loss of generality let $a \in L\left(F^{k(\tau)}\right) \setminus L\left(G^{k(\tau)}\right) = L(F) \setminus L(G)$ (Fact 4-1.). Define $B = A(F) \setminus \{a\}$. Obviously, $(\{a\}, \emptyset, \emptyset) \in \mathcal{L}_\sigma(G \setminus [B, \emptyset]) \setminus \mathcal{L}_\sigma(F \setminus [B, \emptyset])$ showing $F \not\equiv^{\mathcal{L}_\sigma}_{ND} G$ for any semantics $\sigma \in \Phi$. From now on we assume $L\left(F^{k(\tau)}\right) = L\left(G^{k(\tau)}\right)$ and furthermore, $(a, b) \in R\left(F^{k(\tau)}\right) \setminus R\left(G^{k(\tau)}\right)$ for $\tau \in \{stb, ad, co, gr\}$.

1. Assuming $a, b \notin L\left(F^{k(\tau)}\right)$ implies $(a, b) \in R(F) \setminus R(G)$ for $\tau \in \{stb, ad, gr, co\}$ since any τ-kernel require at least one loop for a potential deletion. The following set B will be used throughout the proof and is defined as $B = A(F) \setminus \{a, b\}$. Consider the AFs $F \setminus [B, \emptyset]$ as well as $G \setminus [B, \emptyset]$ abbreviated by $F \setminus B$ or $G \setminus B$, respectively. We list now all possible scenarios.

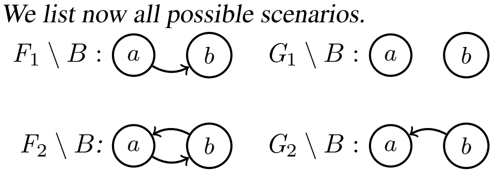

   It can be easily checked that for any $i, j \in \{1, 2\}$, $\mathcal{L}_\sigma(F_i \setminus B) \neq \mathcal{L}_\sigma(G_j \setminus B)$ for any $\sigma \in \Phi$. Consequently, $F \not\equiv^{\mathcal{L}_\sigma}_{ND} G$ for $\sigma \in \Phi$.

2. Assume $a \in L\left(F^{k(\tau)}\right), b \notin L\left(F^{k(\tau)}\right)$.
   (a) In case of $\tau = stb$ we have nothing to show since $(a, b) \in R\left(F^{k(stb)}\right)$ and $a \in L\left(F^{k(stb)}\right)$ is impossible.
   (b) Consider now $\tau = ad$. Let $\sigma \in \{ss, eg, pr, il\}$. In accordance with the admissible kernel we find the following possibilities.

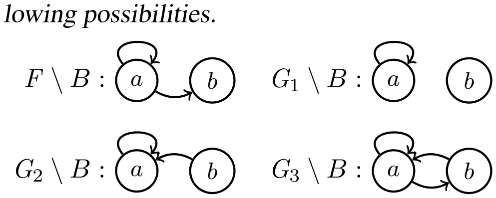

   Obviously, for any $\sigma \in \{ss, eg, pr, il\}$, $\{(\emptyset, \emptyset, \{a, b\})\} = \mathcal{L}_\sigma(F \setminus B)$. Furthermore, $\mathcal{L}_\sigma(G_1 \setminus B) = \{(\{b\}, \emptyset, \{a\})\}$ and $\mathcal{L}_\sigma(G_2 \setminus B) = \mathcal{L}_\sigma(G_3 \setminus B) = \{(\{b\}, \{a\}, \emptyset)\}$. Consequently, $F \not\equiv^{\mathcal{L}_\sigma}_{ND} G$ for $\sigma \in \{ss, eg, pr, il\}$.
   (c) Let $\tau \in \{gr, co\}$. In consideration of the grounded and complete kernel we find the following possibilities.

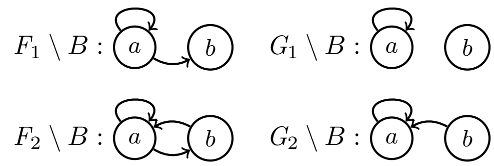

   Let $i \in \{1, 2\}$. We have $\mathcal{L}_{gr}(F_i \setminus B) = \{(\emptyset, \emptyset, \{a, b\})\} \subseteq \mathcal{L}_{co}(F_i \setminus B)$ (Fact 1-1.). On the other hand, for any complete labelling $M \in \mathcal{L}_{co}(G_i \setminus B)$ we necessarily have $M^I(b)$ since it is unattacked in $G_i \setminus B$. Hence, $F \not\equiv^{\mathcal{L}_\sigma}_{ND} G$ for $\sigma \in \{gr, co\}$. Finally, one may easily verify that $(\{b\}, \emptyset, \{a\}) \in \mathcal{L}_{ad}(G_i \setminus B) \setminus \mathcal{L}_{ad}(F_i \setminus B)$. Consequently, $F \not\equiv^{\mathcal{L}_{ad}}_{ND} G$.

3. Assume $a \notin L\left(F^{k(\tau)}\right), b \in L\left(F^{k(\tau)}\right)$.
   (a) We start with $\tau \in \{stb, ad, co\}$. In accordance with the stable, admissible as well as complete kernel we find the following possibilities.

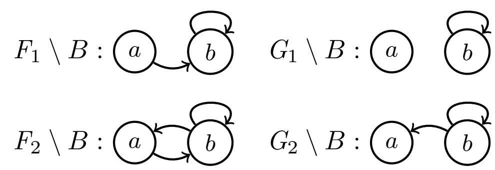

   For any $i \in \{1, 2\}$ and $\sigma \in \{stb, ss, eg, pr, il, ad\}$ we have, $(\{a\}, \{b\}, \emptyset) \in \mathcal{L}_\sigma(F_i \setminus B) \subseteq \mathcal{L}_{co}(F_i \setminus B)$. Obviously, $\mathcal{L}_{stb}(G_i \setminus B) = \emptyset$. Furthermore, for any other complete as well as admissible labelling M of $G_i \setminus B$ we necessarily have $M^U(b)$ since it cannot be in-labelled because it is self-defeating and furthermore, an out-labelling is impossible too because it lacks an attack from an in-labelled argument. Consequently, using Fact1-1. as well as Fact1-3. we derive, $\mathcal{L}_\sigma(F_i \setminus B) \neq \mathcal{L}_\sigma(G_i \setminus B)$ yielding $F \not\equiv^{\mathcal{L}_\sigma}_{ND} G$ for any $\sigma \in \{stb, ss, eg, ad, pr, il, co\}$.
   (b) Consider now $\tau \in \{gr\}$. We find the following possibilities. Note that there is only one possibility for F.

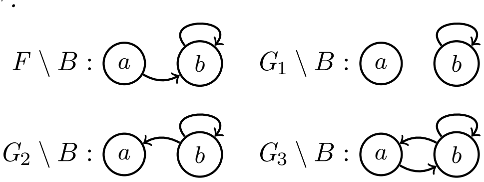

   We have $\mathcal{L}_{gr}(F \setminus B) = \{(\{a\}, \{b\}, \emptyset)\}$. On the other hand, for any grounded labelling $M \in \mathcal{L}_{gr}(G_i \setminus B)$ we find $M^U(b)$ since b is not attacked by an in-labelled argument. This means, $F \not\equiv^{\mathcal{L}_{gr}}_{ND} G$.

4. Assuming $a, b \in L\left(F^{k(\tau)}\right)$ yields a contradiction since for any $\tau \in \{stb, ad, gr, co\}$, $(a, b) \in R\left(F^{k(\tau)}\right)$ is impossible.

Summary and Comparison

The attentive reader may have noticed that we provided only 61 kernel-based characterizations out of 64 defined relations. More precisely, we do not have characterized local expansion equivalence w.r.t. stable, grounded as well as complete labelling-based semantics. Similarly to their extension-based versions it can be checked that none of the existing kernels serve as a characterizing kernel in case of labelling-based semantics. A detailed study will be part of future work.

The following figure provides an overview over all characterization results proven in this paper. Analogously to Figure 2 the entry k in row M and column σ indicates that k characterizes $\equiv_M^{\mathcal{L}_\sigma}$. A red-highlighted entry reflects the situation that extension-based and labelling-based version do not coincide.

	stb	ss	eg	ad	pr	il	gr	co
L	?	$k(ad)$	$k(ad)$	$k(co)$	$k(ad)$	$k(ad)$	?	?
E	$k(stb)$	$k(ad)$	$k(ad)$	$k(co)$	$k(ad)$	$k(ad)$	$k(gr)$	$k(co)$
N	$k(stb)$	$k(ad)$	$k(ad)$	$k(co)$	$k(ad)$	$k(ad)$	$k(gr)$	$k(co)$
S	$k(stb)$	$k(ad)$	$k(ad)$	$k(co)$	$k(ad)$	$k(ad)$	$k(gr)$	$k(co)$
ND	$k(stb)$	$k(ad)$	$k(ad)$	$k(co)$	$k(ad)$	$k(ad)$	$k(gr)$	$k(co)$
D	id	id	id	id	id	id	id	id
LD	id	id	id	id	id	id	id	id
U	id	id	id	id	id	id	id	id

Figure 3: Labelling-based Characterizations

In contrast to extension-based semantics we observe a much more homogeneous picture. Firstly, there is no need for the more sophisticated σ-*-kernels. Secondly, normal deletion equivalence w.r.t. labelling-based semantics is naturally incorporated in the overall picture in the sense that it coincides with its corresponding expansion, normal expansion and strong expansion equivalence notions. Finally, somehow surprisingly, expansion equivalence as well as its local, normal and strong variant w.r.t. admissible labellings are characterizable through the complete (instead of the admissible) kernel.

Conclusions and Related Work

In this paper we studied the problem of characterizing equivalence notions for labelling-based semantics of abstract argumentation frameworks. The labelling-based approach can be traced back to the 1990s (Pollock 1995; Jakobovits and Vermeir 1999). Later on and due to (Caminada 2006) complete labellings were introduced. These labellings play a central role in abstract argumentation, since almost all main-stream extension-based semantics can be reformulated in terms of complete labellings. In recent times the labelling-based approach has received more and more attention. including labelling-based algorithms as well as proof procedures (Caminada 2007; Verheij 2007; Thang, Dung, and Hung 2009; Modgil and Caminada 2009; Caminada 2010; Nofal, Atkinson, and Dunne 2014) and dynamics in argumentation (Boella et al. 2011; Booth et al. 2013; Coste-Marquis et al. 2014; Rienstra, Sakama, and van der Torre 2015). We expect that the results we have obtained will be useful for the aforementioned topics. The main reason for this is that the presented characterization theorems show that there is redundant information even in dynamic scenarios which thus give rise for simplification.

One work which must be mentioned in the context of replacement and labellings is (Baroni et al. 2014). The authors studied the so-called *Input/Output behaviour* of argumentation frameworks. Roughly speaking, the main idea is to consider an argumentation framework as a kind of black box, which receives some input from the external world (i.e, a set of external arguments) via incoming attacks and produces an output to the external world via outgoing attacks. Such an interacting module is called an argumentation multipole. Two multipoles connected with the same external world are considered as *Input/Output equivalent* if the effects, i.e. the produced labellings for external arguments are the same for any reasonable input-labelling. This notion yields the possibility of replacing a multipole with another one embedded in a larger framework without affecting the labellings of the unmodified part of the initial framework. Input/Output equivalence is obviously a less demanding notion than strong equivalence since it requires equal labellings for the external connected part only.

In (Oikarinen and Woltran 2011) the authors said: "Another direction of future work is to consider strong equivalence with respect to labellings rather than with respect to extensions. This would lead to an even more careful notion of equivalence, since labeling – roughly speaking – indicate different reasons why an argument is not contained in an extension. However, since our characterizations based on extensions already provide very strict requirements for the argumentation frameworks under comparison, we do not expect significant further insight from such notions of strong equivalence". As we have seen in Figure 3, there is a majority of equivalence relations where labelling-based and extension-based versions coincide. However, we also observed some exemptions. For instance, the labelling-based versions of normal deletion equivalence behave completely different since they exclude the possibility of possessing different arguments. Furthermore, quite surprisingly, expansion equivalence w.r.t. admissible labellings as well as its local, normal and strong variant are characterizable through the complete kernel. It will be part of future work to study how the labelling versions of conflict-free-based semantics will integrate in the overall picture (Gaggl and Woltran 2011).

References

Baroni, P.; Boella, G.; Cerutti, F.; Giacomin, M.; van der Torre, L. W. N.; and Villata, S. 2014. On the input/output behavior of argumentation frameworks. *Artificial Intelligence* 144–197.

Baroni, P.; Caminada, M.; and Giacomin, M. 2011. An introduction to argumentation semantics. *The Knowledge Engineering Review* 365–410.

Baumann, R., and Brewka, G. 2010. Expanding argumentation frameworks: Enforcing and monotonicity results. In *Computational Models of Argument*, 75–86.

Baumann, R., and Brewka, G. 2015. The equivalence zoo for dung-style semantics. *Journal of Logic and Computation*.

Baumann, R., and Spanring, C. 2015. Infinite argumentation frameworks - on the existence and uniqueness of extensions. In *Advances in Knowledge Representation, Logic Programming, and Abstract Argumentation - Essays Dedicated to Gerhard Brewka on the Occasion of His 60th Birthday*, 281–295.

Baumann, R., and Strass, H. 2015. Open problems in abstract argumentation. In *Advances in Knowledge Representation, Logic Programming, and Abstract Argumentation - Essays Dedicated to Gerhard Brewka on the Occasion of His 60th Birthday*, 325–339.

Baumann, R., and Strass, H. 2016. An abstract logical approach to characterizing strong equivalence in logic-based knowledge representation formalisms. In *16th International Conference on Principles of Knowledge Representation and Reasoning (KR'16) (this volume)*.

Baumann, R., and Woltran, S. 2014. The role of self-attacking arguments in characterizations of equivalence notions (accepted). *Journal of Logic and Computation: Special Issue on Loops in Argumentation*.

Baumann, R. 2012. Normal and strong expansion equivalence for argumentation frameworks. *Artificial Intelligence* 18–44.

Baumann, R. 2014a. Context-free and Context-sensitive kernels: Update and deletion equivalence in abstract argumentation. In *European Conference on Artificial Intelligence*.

Baumann, R. 2014b. *Metalogical Contributions to the Nonmonotonic Theory of Abstract Argumentation*. College Publications - Studies in Logic.

Bench-Capon, T. J. M., and Dunne, P. E. 2007. Argumentation in artificial intelligence. *Artificial Intelligence* 619–641.

Boella, G.; Gabbay, D. M.; Perotti, A.; van der Torre, L.; and Villata, S. 2011. Conditional labelling for abstract argumentation. In *Workshop Theory and Applications of Formal Argumentation*, 232–248.

Booth, R.; Kaci, S.; Rienstra, T.; and van der Torre, L. 2013. A logical theory about dynamics in abstract argumentation. In Liu, W.; Subrahmanian, V.; and Wijsen, J., eds., *Scalable Uncertainty Management*. Springer Berlin Heidelberg. 148–161.

Caminada, M., and Amgoud, L. 2007. On the evaluation of argumentation formalisms. *Artificial Intelligence* 286–310.

Caminada, M., and Gabbay, D. 2009. A logical account of formal argumentation. *Studia Logica* 93:109–145.

Caminada, M., and Pigozzi, G. 2011. On judgment aggregation in abstract argumentation. *Autonomous Agents and Multi-Agent Systems* 22:64–102.

Caminada, M. 2006. On the issue of reinstatement in argumentation. In *10th European Conference on Logics in Artificial Intelligence*, 111–123.

Caminada, M. 2007. An algorithm for computing semi-stable semantics. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, 222–234.

Caminada, M. 2010. An algorithm for stage semantics. In *Computational Models of Argument*, 147–158.

Caminada, M. 2011. A labelling approach for ideal and stage semantics. *Argument & Computation* 2:1–21.

Coste-Marquis, S.; Konieczny, S.; Mailly, J.-G.; and Marquis, P. 2014. On the revision of argumentation systems: Minimal change of arguments statuses. In *14th International Conference on Principles of Knowledge Representation and Reasoning (KR'14)*.

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

Dunne, P. E.; Dvorák, W.; Linsbichler, T.; and Woltran, S. 2015. Characteristics of multiple viewpoints in abstract argumentation. *Artificial Intelligence* 153–178.

Dvorák, W., and Woltran, S. 2014. On the intertranslatability of argumentation semantics. *CoRR*.

Gaggl, S. A., and Woltran, S. 2011. Strong equivalence for argumentation semantics based on conflict-free sets. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, 38–49.

Jakobovits, H., and Vermeir, D. 1999. Robust semantics for argumentation frameworks. *Journal of Logic and Computation* 215–261.

Lifschitz, V.; Pearce, D.; and Valverde, A. 2001. Strongly equivalent logic programs. *ACM Transactions on Computational Logic* 526–541.

Modgil, S., and Caminada, M. W. 2009. Proof theories and algorithms for abstract argumentation frameworks. In Rahwan, I., and Simari, G., eds., *Argumentation in Artificial Intelligence*. Springer Publishing Company, Incorporated. 105–129.

Nofal, S.; Atkinson, K.; and Dunne, P. E. 2014. Algorithms for decision problems in argument systems under preferred semantics. *Artificial Intelligence* 23–51.

Oikarinen, E., and Woltran, S. 2011. Characterizing strong equivalence for argumentation frameworks. *Artificial Intelligence* 1985–2009.

Pollock, J. L. 1995. *Cognitive Carpentry: A Blueprint for How to Build a Person*. Cambridge, Massachusetts: The MIT Press.

Rahwan, I., and Simari, G. R., eds. 2009. *Argumentation in Artificial Intelligence*. Springer.

Rienstra, T.; Sakama, C.; and van der Torre, L. 2015. Persistence and monotony properties of argumentation semantics. In *Workshop Theory and Applications of Formal Argumentation*.

Thang, P. M.; Dung, P. M.; and Hung, N. D. 2009. Towards a common framework for dialectical proof procedures in abstract argumentation. *Journal of Logic and Computation* 1071–1109.

Truszczynski, M. 2006. Strong and uniform equivalence of nonmonotonic theories - an algebraic approach. *Annals of Mathematics and Artificial Intelligence* 245–265.

Turner, H. 2001. Strong equivalence for logic programs and default theories (made easy). In *Logic Programming and Nonmonotonic Reasoning LPNMR*, 81–92.

Turner, H. 2004. Strong equivalence for causal theories. In *International Conference on Logic Programming and Nonmonotonic Reasoning*, 289–301.

Verheij, B. 2007. A labeling approach to the computation of credulous acceptance in argumentation. In *20th International Joint Conference on Artificial Intelligence IJCAI*, 623–628.