

Monotonicity and Noise-Tolerance in Case-Based Reasoning with Abstract Argumentation

Guilherme Paulino-Passos, Francesca Toni

Department of Computing, Imperial College London

{g.passos18, f.toni}@imperial.ac.uk

Abstract

Recently, abstract argumentation-based models of case-based reasoning (*AA-CBR* in short) have been proposed, originally inspired by the legal domain, but also applicable as classifiers in different scenarios. However, the formal properties of *AA-CBR* as a reasoning system remain largely unexplored. In this paper, we focus on analysing the non-monotonicity properties of a *regular* version of *AA-CBR* (that we call *AA-CBR_≥*). Specifically, we prove that *AA-CBR_≥* is not cautiously monotonic, a property frequently considered desirable in the literature. We then define a variation of *AA-CBR_≥* which is cautiously monotonic. Further, we prove that such variation is equivalent to using *AA-CBR_≥* with a restricted casebase consisting of all “surprising” and “sufficient” cases in the original casebase. As a by-product, we prove that this variation of *AA-CBR_≥* is cumulative, rationally monotonic, and empowers a principled treatment of noise in “incoherent” casebases. Finally, we illustrate *AA-CBR* and cautious monotonicity questions on a case study on the U.S. Trade Secrets domain, a legal casebase.

1 Introduction

Case-based reasoning (CBR) relies upon known solutions for problems (past cases) to infer solutions for unseen problems (new cases), based upon retrieving past cases which are “similar” to the new cases. It is widely used in legal settings (e.g. Prakken et al. 2015; Čyras, Satoh, and Toni 2016a), for classification (e.g. via the k-NN algorithm and, recently, within the DEAr methodology (Cocarascu et al. 2020)) and for explanation (e.g. see Nugent and Cunningham 2005; Kenny and Keane 2019; Cocarascu et al. 2020).

In this paper we focus on a recent approach to CBR based upon an argumentative reading of (past and new) cases (Čyras, Satoh, and Toni 2016a; Čyras, Satoh, and Toni 2016b; Cocarascu, Čyras, and Toni 2018; Čyras et al. 2019; Cocarascu et al. 2020), and using *Abstract Argumentation (AA)* (Dung 1995) as the underpinning machinery. We will refer to all proposed incarnations of this approach in the literature generically as *AA-CBR* (the acronym used in the original paper (Čyras, Satoh, and Toni 2016a)): they all generate an AA framework from a CBR problem – a graph structure where cases are arguments, “more specific” past cases attack “less specific” past cases or a “default argument” (which embeds a sort of bias), and new cases at-

tack “irrelevant” past cases; then, CBR is reduced to testing membership of this default argument in the grounded extension (Dung 1995). The use of argumentation in *AA-CBR* naturally paves the way towards explanation generation for CBR tasks, e.g. in the form of dispute trees (Čyras, Satoh, and Toni 2016b; Cocarascu et al. 2020) or excess features (Čyras et al. 2019), possibly for supporting interactions with users, in the spirit of recent research on incorporating feedback in recommender systems (Rago et al. 2021) and showing influence structures from neural network classifiers (Dejl et al. 2021).

Different incarnations of *AA-CBR* use different mechanisms for defining the aforementioned “specificity”, “irrelevance” and “default argument”: the original version in (Čyras, Satoh, and Toni 2016a) is applicable only to cases characterised by sets of features and defines all three notions in terms of subsets, while the version used for classification in (Cocarascu et al. 2020) defines specificity in terms of a generic partial order, irrelevance in terms of a generic relation and default argument in terms of a generic characterisation. Thus, it is in principle applicable to cases characterised in any way, as sets of features or unstructured (Cocarascu et al. 2020). We will study a special, *regular* instance (*AA-CBR_≥*) of this more recent presentation, in which irrelevance and the default argument are both defined via specificity (and in particular the default argument is defined in terms of the most specific case). *AA-CBR_≥* admits the original *AA-CBR* in (Čyras, Satoh, and Toni 2016a) as an instance, obtained by choosing the partial order to be the subset relation. Differently from previous presentations, *AA-CBR_≥* accepts “incoherent” casebases, in which there is “noise” in the form of different cases with the same characterisation but different outcomes. Incoherence may result from a limited language to express features or genuine errors in generating the casebases. Independently of the reasons behind incoherence, and especially when this is outside the control of reasoning system designers, it is important that the reasoning system is able to tolerate it.

AA-CBR was originally inspired by the legal domain in (Čyras, Satoh, and Toni 2016a), but some incarnations of *AA-CBR*, integrating dynamic features, have proven useful in predicting and explaining the passage of bills in the UK Parliament (Čyras et al. 2019), and instantiations of the more

generic version of Cocarascu et al. (2020) have shown to be fruitfully applicable as classifiers (Cocarascu et al. 2020). We study *non-monotonicity* properties of $AA-CBR_{\succeq}$ understood at the same time as a reasoning system and as a classifier. These properties, typically considered in logic, intuitively characterise in which sense systems may stop inferring some conclusions when more information becomes available (Makinson 1994). These properties are thus related to modelling inference which is tentative and defeasible, as opposed to the indefeasible form of inference of classical logic. Non-monotonicity properties have already been studied in argumentation, e.g. for ABA, ABA+ (Čyras and Toni 2015; Čyras and Toni 2016), $ASPIC^+$ (Dung 2014; Dung 2016) and logic-based argumentation (Hunter 2010). We study them for the application of argumentation to classification via $AA-CBR_{\succeq}$.

Specifically, we prove that the kind of inference underpinning $AA-CBR_{\succeq}$ lacks a standard non-monotonicity property, namely *cautious monotonicity*, sanctioning, intuitively, that if a conclusion is added to the set of premises (here, the casebase), then no conclusion is lost, that is, everything which was inferable still is so. In terms of a supervised classifier, satisfying cautious monotonicity culminates in being “closed” under self-training. That is, augmenting the dataset with conclusions inferred by the classifier itself does not change the classifier. Then, we make a two-fold contribution: we define (formally and algorithmically) a provably cautiously monotonic variant of $AA-CBR_{\succeq}$, that we call $cAA-CBR_{\succeq}$, and prove that it is equivalent to $AA-CBR_{\succeq}$ applied to a restricted casebase consisting of all “surprising” and “sufficient” cases in the original casebase. We also show that cautious monotonicity of $cAA-CBR_{\succeq}$ leads to the desirable properties of *cumulativity* and *rational monotonicity*, and that, as a by-product, our cautiously monotonic variant leads to a desirable treatment of noise in incoherent casebases. It is interesting that $cAA-CBR_{\succeq}$ deals with it serendipitously, as a direct consequence of insuring cautious monotonicity.

This paper generalises our previous work (Paulino-Passos and Toni 2020) by also dealing with incoherent casebases, presenting a case study, and discussing the position of our contribution in the related literature. We omit some proofs and details for lack of space but they are available in the appendix at <https://arxiv.org/abs/2107.06413>.

2 Motivating Illustration

In this section we introduce a simple setting for the informal illustration of the original $AA-CBR$, its non-monotonicity and the desirability of some restrictions thereof, as well as problems raised by the presence of incoherence in the case base when deploying $AA-CBR$. Thus, this section serves as a motivating illustration for our approach, which restricts non-monotonicity and is incoherence-tolerant.

Example 1 (Non-Monotonicity). Consider a simplified legal system built by cases and adhering, like most modern legal systems, to the principle by which, unless proven otherwise, no person is to be considered guilty of a crime. This can be represented by a “default argument” $(\emptyset, -)$, indicat-

ing that, in the absence of any information about any person, the legal system should infer a negative outcome $-$ (that the person is *not* guilty). $(\emptyset, -)$ can be understood as an argument, in the AA sense, given that it is merely what is called a relative presumption, since it is open to proof to the contrary, e.g. by proving that the person did indeed commit a crime. Let us consider here one possible crime: homicide¹ (hm). In one case, it was established that the defendant committed homicide, and he was considered guilty, represented as $(\{hm\}, +)$.

Consider now a new case $(\{hm, sd\}, ?)$, with an unknown outcome, of a defendant who committed homicide, but for which it was proven that it was in self-defence (sd). In order to predict the new case’s outcome by CBR, $AA-CBR$ reduces the prediction problem to that of membership of the default argument in the grounded extension (Dung 1995) \mathbb{G} of the AA framework in Figure 1a: given that $(\emptyset, -) \notin \mathbb{G}$, the predicted outcome is positive (i.e. guilty), disregarding sd and, indeed, no matter what other feature this case may have. Thus, up to this point, having the feature hm is a sufficient condition for predicting guilty. If, however, the courts decide that for this new case the defendant should be acquitted, the case $(\{hm, sd\}, -)$ enters in our casebase. Now, having the feature hm is no longer a sufficient condition for predicting guilty, and any case with both hm and sd will be predicted a negative outcome (i.e. that the person is innocent). This is the case for predicting the outcome of a new case with again both hm and sd , in $AA-CBR$ using the AA framework in Figure 1b. Thus, adding a new case to the casebase removed some conclusions which were inferred from the previous, smaller casebase, showing that $AA-CBR$ is indeed non-monotonic. This does not mean that some restrictions on non-monotonicity might not be desirable. For instance, we might expect in a legal system that, if for the current case law, two cases are to be judged in a certain way, then one of the cases happening in court and indeed being decided in that way would not affect the body of case law itself, thus the outcome for the second case would be expected to be unchanged.

The following example illustrates the challenges posed by incoherent casebases, with noisy cases, in $AA-CBR$.

Example 2 (Noise-intolerance). Consider a different augmentation of the initial casebase in Example 1, resulting in the casebase $\{(\{hm\}, +), (\{hm\}, -)\}$, whereby a positive and a negative outcomes are recorded for exactly the same profile for defendants. This can be deemed to be “incoherent”². In $AA-CBR$, this case would have a positive outcome (guilty), since the default argument is not in the grounded extension. However, this seems unsatisfactory, since the default argument is not attacked by any argument in

¹This is a toy example, so the terms used do not correspond to a specific jurisdiction.

²This casebase may result from a limited language for characterising cases, e.g. ignoring the possibility of indicating core differences between defendants, such as that the defendant characterised by $(\{hm\}, -)$ is a minor. We assume here that this incoherence cannot be rectified by a language variation, or simply that it comes from the data, and we cannot remove it based on the data alone.

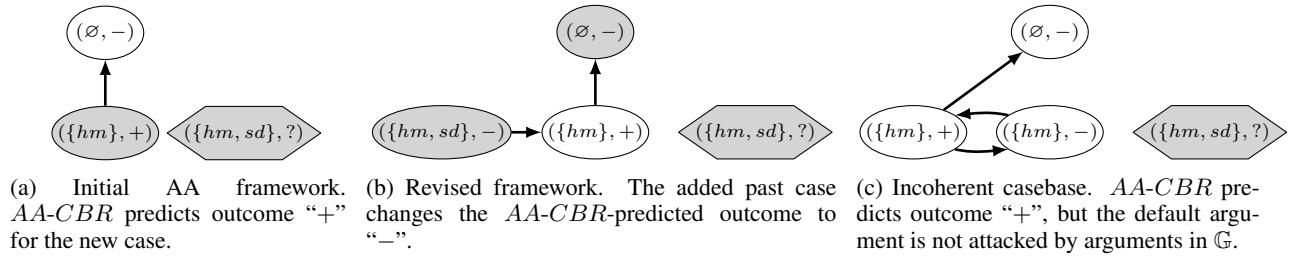


Figure 1: AA frameworks when using $AA-CBR$ for Examples 1 and 2. Past cases (with outcomes) and new case (with unknown outcome) are arguments. (Grounded extensions \mathbb{G} are shaded.)

the grounded extension of the corresponding AA framework (see Figure 1c). In some sense, this means that no past case serves as ground for the decision, but the outcome is guilty nonetheless. Thus, what is deciding the outcome is not a relevant case per se, but the incoherence itself. This can be regarded as a form of intolerance to noise. Note that incoherence does not always give this form of noise intolerance in $AA-CBR$: if a past case $(\{hm, sd\}, +)$ were included, the outcome would be the same, but the default argument would be attacked by (some argument in) the grounded extension.

Our form of $AA-CBR$ is guaranteed to be tolerant of noise always, independently of the casebase.

3 Preliminaries

Abstract Argumentation. An *abstract argumentation framework* (AF) (Dung 1995) is a pair $(Args, \rightsquigarrow)$, where $Args$ is a set (of *arguments*) and $\rightsquigarrow \subseteq Args \times Args$ is a binary relation on $Args$. For $\alpha, \beta \in Args$, if $\alpha \rightsquigarrow \beta$, then we say that α *attacks* β and that α is an *attacker* of β . For a set of arguments $E \subseteq Args$ and an argument $\alpha \in Args$, E *defends* α if for all $\beta \rightsquigarrow \alpha$ there exists $\gamma \in E$ such that $\gamma \rightsquigarrow \beta$. Then, the *grounded extension* of $(Args, \rightsquigarrow)$ can be constructed as $\mathbb{G} = \bigcup_{i \geq 0} G_i$, where G_0 is the set of all unattacked arguments, and $\forall i \geq 0$, G_{i+1} is the set of arguments that G_i defends. For any $(Args, \rightsquigarrow)$, the grounded extension \mathbb{G} always exists and is unique and, if $(Args, \rightsquigarrow)$ is well-founded (Dung 1995), extensions under other semantics (e.g. stable extensions) are equal to \mathbb{G} . For finite AFs, $(Args, \rightsquigarrow)$ is well-founded iff it is acyclic. Given $(Args, \rightsquigarrow)$, we will sometimes use $\alpha \in (Args, \rightsquigarrow)$ to stand for $\alpha \in Args$.

Non-monotonicity Properties. We will be interested in the following properties.³ An arbitrary inference relation \vdash (for a language including, in particular, sentences a, b , etc., with negations $\neg a$ and $\neg b$, etc., and sets of sentences A, B) is said to satisfy:

1. *non-monotonicity*, iff $A \vdash a$ and $A \subseteq B$ do not imply that $B \vdash a$;
2. *cautious monotonicity*, iff $A \vdash a$ and $A \vdash b$ imply that $A \cup \{a\} \vdash b$;
3. *cut*, iff $A \vdash a$ and $A \cup \{a\} \vdash b$ imply that $A \vdash b$;

³We are mostly following the treatment of (Makinson 1994).

4. *cumulativity*, iff \vdash is both cautiously monotonic and satisfies cut;
5. *rational monotonicity*, iff $A \vdash a$ and $A \not\vdash \neg b$ imply that $A \cup \{b\} \vdash a$;
6. *completeness*, iff either $A \vdash a$ or $A \vdash \neg a$.

4 Abstract Argumentation for Case-Based Reasoning

Here, we define $AA-CBR_{\succeq}$, adapting definitions from (Cocarascu et al. 2020). All incarnations of $AA-CBR$, including $AA-CBR_{\succeq}$, map a *dataset* D of *examples* labelled with an *outcome* and an *unlabelled example* (with unknown outcome) into an AF. The dataset may be understood as a *casebase*, the labelled examples as *past cases* and the unlabelled example as a *new case*: we will use these terminologies interchangeably throughout. In this paper, as in (Cocarascu et al. 2020), examples/cases have a characterisation (e.g., as in (Čyras, Satoh, and Toni 2016a), characterisations may be sets of features), and outcomes are chosen from two available ones, one of which is selected up-front as the *default outcome*. Finally, in the spirit of (Cocarascu et al. 2020), we assume that the set of characterisations of (past and new) cases is equipped with a partial order \preceq (whereby $\alpha \prec \beta$ holds if $\alpha \preceq \beta$ and $\alpha \neq \beta$ and is read “ α is less *specific* than β ”) and with a relation $\not\prec$ (whereby $\alpha \not\prec \beta$ is read as “ β is *irrelevant* to α ”). Formally:

Definition 3 (Adapted from (Cocarascu et al. 2020)). Let X be a set of *characterisations*, equipped with partial order \prec and binary relation $\not\prec$. Let $Y = \{\delta_o, \delta_o\}$ be the set of (all possible) *outcomes*, with δ_o the *default outcome*. Then, a *casebase* D is a finite set such that $D \subseteq X \times Y$ (thus a *past case* $\alpha \in D$ is of the form (α_C, α_o) for $\alpha_C \in X$, $\alpha_o \in Y$) and a *new case* is of the form $(N_C, ?)$ for $N_C \in X$. We also discriminate a particular element $\delta_C \in X$ and define the *default argument* $(\delta_C, \delta_o) \in X \times Y$.

A casebase D is *coherent* iff there are no two cases $(\alpha_C, \alpha_o), (\beta_C, \beta_o) \in D$ such that $\alpha_C = \beta_C$ but $\alpha_o \neq \beta_o$, and it is *incoherent* otherwise.

For simplicity of notation, we sometimes extend the definition of \succeq to $X \times Y$, by setting $(\alpha_c, \alpha_o) \succeq (\beta_c, \beta_o)$ iff $\alpha_c \succeq \beta_c$.⁴

⁴In (Cocarascu et al. 2020), \succeq was directly given over $X \times Y$. Note that, when D is coherent, our “lifted” \succeq is guaranteed to be a

Definition 4 (Adapted from (Cocarascu et al. 2020)). The AF mined from a dataset D and a new case $(N_C, ?)$ is $(Args, \rightsquigarrow)$, in which:

- $Args = D \cup \{(\delta_C, \delta_o)\} \cup \{(N_C, ?)\}$;
- for $(\alpha_C, \alpha_o), (\beta_C, \beta_o) \in D \cup \{(\delta_C, \delta_o)\}$, it holds that $(\alpha_C, \alpha_o) \rightsquigarrow (\beta_C, \beta_o)$ iff
 1. $\alpha_o \neq \beta_o$,
 2. $\alpha_C \succeq \beta_C$, and
 3. $\nexists (\gamma_C, \gamma_o) \in D \cup \{(\delta_C, \delta_o)\}$ with $\alpha_C \succ \gamma_C \succ \beta_C$ and $\gamma_o = \alpha_o$;
- for $(\beta_C, \beta_o) \in D \cup \{(\delta_C, \delta_o)\}$, it holds that $(N_C, ?) \rightsquigarrow (\beta_C, \beta_o)$ iff $(N_C, ?) \not\rightsquigarrow (\beta_C, \beta_o)$.

The AF mined from a dataset D alone is $(Args', \rightsquigarrow')$, with $Args' = Args \setminus \{(N_C, ?)\}$ and $\rightsquigarrow' = \rightsquigarrow \cap (Args' \times Args')$.

Note that if D is coherent, then the “equals” case in item 2 of the definition of attack will never apply. As a result, the AF mined from a coherent D (and any $(N_C, ?)$) is guaranteed to be well-founded, in the sense of Dung (1995).

Definition 5 (Adapted from (Cocarascu et al. 2020)). Let \mathbb{G} be the grounded extension of the AF mined from D and $(N_C, ?)$, with default argument (δ_C, δ_o) . The outcome for N_C is δ_o if (δ_C, δ_o) is in \mathbb{G} , and $\bar{\delta}_o$ otherwise.

In this paper we focus on a particular case of this scenario:

Definition 6. The AF mined from D alone and the AF mined from D and $(N_C, ?)$, with default argument (δ_C, δ_o) , are *regular* when the following holds:

1. the irrelevance relation $\not\sim$ is defined as: $x_1 \not\sim x_2$ iff $x_1 \not\preceq x_2$, and
2. δ_C is the least element of X .⁵

This restriction connects the treatment of a characterisation α_C as a new case and as a past case and is necessary in order to satisfy desirable properties, such as a relation between new cases and “nearest” past cases, which is proven to hold in the appendix⁶.

From now on, we will restrict our attention to regular mined AFs. We will refer to the (regular) AF mined from D and $(N_C, ?)$, with default argument (δ_C, δ_o) , as $AF_{\succeq}(D, N_C)$, and to the (regular) AF mined from D alone as $\bar{AF}_{\succeq}(D)$. Also, for short, given $AF_{\succeq}(D, N_C)$, with default argument (δ_C, δ_o) , we will refer to the outcome for N_C as $AA-CBR_{\succeq}(D, N_C)$.⁷ Unless otherwise stated, we will assume arbitrary $X, Y, D, (N_C, ?)$, and (δ_C, δ_o) (satisfying the previously defined constraints). Finally, we will refer to $AA-CBR_{\succeq}$ instantiated with $\succeq = \supseteq$ and $(\delta_C, \delta_o) = (\emptyset, -)$ as $AA-CBR_{\supseteq}$.

partial order on $X \times Y$ (and thus equivalent to the one in (Cocarascu et al. 2020)), but when D is incoherent anti-symmetry may fail for two cases with different outcomes but same characterisation, and thus \succeq is merely a preorder on $X \times Y$.

⁵Indeed this is not a strong condition, since it can be proved that if $\alpha_C \not\preceq \delta_C$ then all cases (α_C, α_o) in the casebase could be removed, as they would never change an outcome. On the other hand, assuming also the first condition in Definition 6, if $(\alpha_C, ?)$ is the new case and $\alpha_C \not\preceq \delta_C$, then the outcome is $\bar{\delta}_o$ necessarily.

⁶<https://arxiv.org/abs/2107.06413>

⁷In the notation we omit (δ_C, δ_o) , and leave it implicit instead for readability.

5 Non-Monotonicity Analysis of Classifiers

In this section we provide a generic analysis of the non-monotonicity properties of data-driven classifiers, using D, X and Y to denote generic inputs and outputs of classifiers, admitting our casebases, characterisations and outcomes as special instances. Later in the paper, we will apply this analysis to $AA-CBR_{\succeq}$ and our modification thereof. Typically, a classifier can be understood as a function from an input set X to an output set Y . In machine learning, classifiers are obtained by training with an initial, finite $D \subseteq X \times Y$, called the training set. In (any form of) $AA-CBR$, D can also be seen as a training set of sorts. Thus, we will characterise a classifier as a two-argument function \mathbb{C} that maps from a dataset $D \subseteq X \times Y$ and from a new input $x \in X$ to a prediction $y \in Y$.⁸ Notice that this function is total, in line with the common assumptions that classifiers generalise beyond their training dataset.

Let us model directly the relationship between the dataset D and the predictions it makes via the classifier as an inference system in the following way:

Definition 7. Given a classifier $\mathbb{C}: 2^{X \times Y} \times X \rightarrow Y$, let $\mathcal{L} = \mathcal{L}^+ \cup \mathcal{L}^-$ be a language consisting of atoms $\mathcal{L}^+ = X \times Y$ and negative sentences $\mathcal{L}^- = \{\neg(x, y) \mid (x, y) \in X \times Y\}$. Then, $\vdash_{\mathbb{C}}$ is an *inference relation* from $2^{\mathcal{L}^+}$ to \mathcal{L} such that

- $D \vdash_{\mathbb{C}} (x, y)$ iff $\mathbb{C}(D, x) = y$;
- $D \vdash_{\mathbb{C}} \neg(x, y)$ iff there is a y' such that $\mathbb{C}(D, x) = y'$ and $y' \neq y$.⁹

Intuitively, \mathbb{C} defines a language consisting of atoms (representing labelled examples) and their negations, and $\vdash_{\mathbb{C}}$ applies a sort of closed world assumption around \mathbb{C} .

Then, we can study non-monotonicity properties (see Section 3) of $\vdash_{\mathbb{C}}$. Here, completeness (the first item in the theorem) causes them to collapse.

Theorem 8. 1. $\vdash_{\mathbb{C}}$ is complete, i.e. for every $(x, y) \in (X \times Y)$, either $D \vdash_{\mathbb{C}} (x, y)$ or $D \vdash_{\mathbb{C}} \neg(x, y)$.

2. $\vdash_{\mathbb{C}}$ is consistent, i.e. for every $(x, y) \in (X \times Y)$, it does not hold that both $D \vdash_{\mathbb{C}} (x, y)$ and $D \vdash_{\mathbb{C}} \neg(x, y)$.
3. $\vdash_{\mathbb{C}}$ is cautiously monotonic iff it satisfies cut.
4. $\vdash_{\mathbb{C}}$ is cautiously monotonic iff it is cumulative.
5. $\vdash_{\mathbb{C}}$ is cautiously monotonic iff it satisfies rational monotonicity.

⁸Notice that this understanding relies upon the assumption that classifiers are deterministic. Of course this is not the case for many machine learning models, e.g. artificial neural networks trained using stochastic gradient descent and randomised hyperparameter search. This understanding is however in line with recent work using decision functions as approximations of classifiers whose output needs explaining (e.g. see (Shih, Choi, and Darwiche 2019)). Moreover, it works well when analysing $AA-CBR_{\succeq}$.

⁹We could equivalently have defined $D \vdash_{\mathbb{C}} \neg(x, y)$ iff $\mathbb{C}(D, x) \neq y$. We have not done so as the used definition can be generalised for a scenario in which \mathbb{C} is not necessarily a total function. This scenario is left for future work.

6 Cautious Monotonicity in $AA-CBR_{\succeq}$

Our first main result is about (lack of) cautious monotonicity of the inference relation drawn from the classifier $AA-CBR_{\succeq}(D, N_C)$.

Theorem 9. $\vdash_{AA-CBR_{\succeq}}$ is not cautiously monotonic.

Proof. We show a counterexample, choosing $X = 2^{\{a,b,c,z\}}$, $Y = \{-, +\}$, and $\preceq = \supseteq$. Define $D = \{(\{a\}, +), (\{c\}, +), (\{a, b\}, -), (\{c, z\}, -)\}$ and $(\delta_C, \delta_o) = (\emptyset, -)$, and two new cases: $N_1 = \{a, b, c\}$ and $N_2 = \{a, b, c, z\}$.

Consider now $AA-CBR_{\succeq}(D, N_1)$ and $AA-CBR_{\succeq}(D, N_2)$. We can see in Figure 2a that $D \vdash_{AA-CBR_{\succeq}} (N_1, +)$ and in Figure 2b that $D \vdash_{AA-CBR_{\succeq}} (N_2, -)$.

Finally, let us consider $AF_{\succeq}(D \cup \{(N_1, +)\}, N_2)$ in Figure 2c. We can then conclude that $D \cup \{(N_1, +)\} \vdash_{AA-CBR_{\succeq}} (N_2, +)$ even though $D \vdash_{AA-CBR_{\succeq}} (N_1, +)$ and $D \vdash_{AA-CBR_{\succeq}} (N_2, -)$, as required. \square

Note that the proof of Theorem 9 shows that the inference relation drawn from the original $AA-CBR$ (i.e. $AA-CBR_{\supseteq}$) is also non-cautiously monotonic, given that the proof's counterexample is obtained with $AA-CBR_{\supseteq}$. Also, note that the counterexample amounts to an expansion of Example 1, as follows.

Example 10. (Example 1 cont.) Assume that a different type of crime happened: public offending someone's honour, which we will call defamation (*df*). In one case, it was established that the defendant did publicly damage someone's honour, and was considered guilty ($\{df\}, +$). In a subsequent case, even if proven that the defendant did hurt someone's honour, it was established that this was done by a true allegation, and thus the case was dismissed, represented as $(\{df, td\}, -)$. What happens, then, if a same defendant is: 1. simultaneously proven guilty of homicide, of defamation, but shown to have committed the homicide in self-defence ($(\{hm, df, sd\}, ?)$); or 2. simultaneously proven guilty of homicide, of defamation, shown to have committed the homicide in self-defence, and also shown to have committed defamation by a true allegation ($(\{hm, df, sd, td\}, ?)$)?

We can map these situations to our counterexample in the proof of Theorem 9 by setting $a = hm$, $b = sd$, $c = df$, and $z = td$. The first question is answered by the AF represented in Figure 2a, with outcome $+$, that is, the defendant is considered guilty. The proof of Theorem 9 shows that the answer to the second question in $AA-CBR_{\succeq}$ would depend on whether the case in the first question was already judged or not. If not, then the cases $(\{hm, sd\}, -)$ and $(\{df, td\}, -)$ would be the nearest cases, and the outcome would be $-$, that is, not guilty. However, if the case in the first question was already judged and incorporated into the case law, it would serve as a counterargument for $(\{hm, sd\}, -)$, and guarantee that the outcome is $+$ (guilty). Intuitively, this seems strange, in particular, because the order in which the case in the first answer is judged affects the case in the second question.

This example aims only to illustrate an interpretation in which $AA-CBR_{\succeq}$ operates seemingly inappropriately.

Whether this behaviour is desirable in general depends on the intended application and other elements such as the relation between the characterisations and the partial order.

7 A Cumulative $AA-CBR_{\succeq}$

We will now present $cAA-CBR_{\succeq}$, a novel, *cumulative* incarnation of $AA-CBR$ satisfying cautious monotonicity.

Concise Dataset. Firstly, let us present some general notions, defined in terms of the \vdash_C inference relation from an arbitrary classifier C . Intuitively, we are after a relation \vdash'_C such that if $D \vdash_C c$ and $D \vdash_C d$, then $D \cup \{c\} \vdash'_C d$ (in our concrete setting, $\vdash_C = \vdash_{AA-CBR_{\succeq}}$ and $\vdash'_C = \vdash_{cAA-CBR_{\succeq}}$). We also want the property that, whenever D is "well-behaved" (in a sense to be made precise later), $D \vdash_C s$ iff $D \vdash'_C s$. In this way, given that $D \vdash'_C c$ and $D \vdash'_C d$, then we would conclude $D \cup \{c\} \vdash'_C d$, making \vdash'_C a cautiously monotonic relation. We will define \vdash'_C by building a subset of the original dataset in such a way that cautious monotonicity is preserved. We start with the following notion of *includable* examples:

Definition 11. An example $(x, y) \in X \times Y$ is *surprising* w.r.t. D iff $D \setminus \{(x, y)\} \not\vdash_C (x, y)$ and *sufficient* w.r.t. D iff $D \cup \{(x, y)\} \vdash_C (x, y)$. Additionally, an example $(x, y) \in X \times Y$ is *includable* w.r.t. D iff it is both surprising and sufficient w.r.t. to D .

The definition of includable example has two parts: that the example is surprising, in the sense that, without it, the predicted outcome would be different, and that it is sufficient, in the sense that adding it makes it inferable. We then define the notion of *concise* subsets, amounting to includable examples only:

Definition 12. Let $S \subseteq X \times Y$ be a dataset, $S' \subseteq S$, and let $\varphi(S') = \{(x, y) \in S' \mid (x, y) \text{ is includable w.r.t. } S'\}$. Then S' is *concise* w.r.t. S whenever it is a fixed point of φ , that is, $\varphi(S') = S'$.

To illustrate in the context of $AA-CBR$, consider S from which the AF in Figure 2c is drawn. S is not concise w.r.t. itself, since $(\{a, b, c\}, +)$ is not includable w.r.t. S (indeed, $S \setminus \{(\{a, b, c\}, +)\} \vdash_{AA-CBR_{\succeq}} (\{a, b, c\}, +)$, see Figure 2a). Also, $S' = S \setminus \{(\{a, b\}, -), (\{a, b, c\}, +)\}$ is not concise either (w.r.t. S), as $(\{a, b\}, -)$ is includable w.r.t. S' (the predicted outcome being $+$), but not an element of S' . The only concise subset of S here is $S'' = S \setminus \{(\{a, b, c\}, +)\}$.

Let us now consider $D' \subseteq D$, for a dataset D . If D' is concise w.r.t. D , $(x, y) \in (X \times Y) \setminus D$ is an example not in D already and $D' \vdash_C (x, y)$, then (x, y) is not includable w.r.t. D' , and thus D' is still concise w.r.t. $D \cup \{(x, y)\}$. Now, suppose that there is exactly one such concise $D' \subseteq D$ w.r.t. D (let us refer to this subset simply as *concise*(D)). Then, it seems attractive to define \vdash'_C as: $D \vdash'_C (x, y)$ iff *concise*(D) $\vdash_C (x, y)$. Such \vdash'_C inference relation would then be cautiously monotonic if *concise*(D) = *concise*($D \cup \{(x, y)\}$). To see that, consider $(x', y') \in X \times Y$ such that $D \vdash'_C (x', y')$. Then, since *concise*(D) = *concise*($D \cup \{(x, y)\}$), for our new \vdash'_C , D

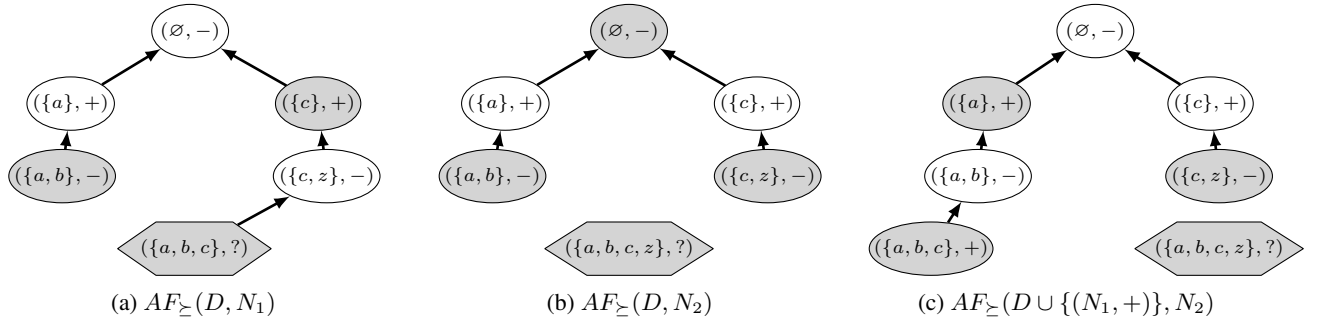


Figure 2: AFs for the proof of Theorem 9, with the grounded extension shaded.

Algorithm 1: *simple_add* for $AA-CBR_{\succeq}$.

Input: An $AA-CBR_{\succeq}$ framework $(Args, \rightsquigarrow)$ and a case $n = (n_C, n_o)$
Output: A new $AA-CBR_{\succeq}$ framework $(Args', \rightsquigarrow')$
 $DEF \leftarrow \{(x, y) \in AF_{\succeq}(Args, n_C) \mid$
 $(x, y) \neq (n_C, ?) \text{ and } (n_C, ?) \text{ defends}$
 $(x, y) \text{ in } AF_{\succeq}(Args, n_C)\};$
 $Args' \leftarrow Args \cup \{n\};$
 $\rightsquigarrow' \leftarrow (\rightsquigarrow \cup \{(n, a) \mid a = (a_C, a_o), a \in DEF,$
 $\text{and } a_o \neq n_o\});$
return $(Args', \rightsquigarrow')$

and $D \cup \{(x, y)\}$ would infer the exact same sentences, thus $D \cup \{(x, y)\} \vdash_C (x', y')$. This equality is indeed guaranteed given that $(x, y) \notin D$, thus it is not includable, and then a concise subset of D is still a concise subset of $D \cup \{(x, y)\}$ (otherwise, if $(x, y) \in D$, the equality would be trivial). Note that concision is too strong a property here: all that is needed is that a subset D' is selected such that every case in it is surprising w.r.t. D' itself. However, concision implies that as many cases are added as possible, while restricting to the ones that guarantee their outcomes.

In the remainder, we state uniqueness and give an algorithm that constructs $concise(D)$, in the case of $AA-CBR_{\succeq}$. If D is incoherent, there might be no concise subset thereof, but our method will still be useful, as we discuss later.

Uniqueness and Algorithm. We first give a property of concise subsets:

Theorem 13. *For $AA-CBR_{\succeq}$, if there is a concise $D' \subseteq D$ w.r.t. D then every concise subset of D w.r.t. D is the same as D' .*

The procedure for finding this unique $concise(D)$, if it exists, is integrated within Algorithm 2, using in turn Algorithm 1. If $concise(D)$ does not exist, the algorithm will still return some $D' \subseteq D$ consisting only of surprising examples w.r.t. D' (in fact, Algorithm 2 returns both this subset and the AF mined from it, for ease of use). The main idea behind the algorithm is simple: we start with the default argument, and progressively build the AF by adding cases from D by following the partial order \preceq . Before adding a past

Algorithm 2: Setup/learning for $cAA-CBR_{\succeq}$.

Input: A dataset D
Output: A subset D' of D , an AF $cAA-CBR_{\succeq}(D)$
 $unprocessed \leftarrow D;$
 $Args \leftarrow \{(\delta_C, \delta_o)\};$
 $\rightsquigarrow \leftarrow \emptyset;$
while $unprocessed \neq \emptyset$ **do**
 $\quad stratum \leftarrow \{(x, y) \in unprocessed \mid$
 $\quad (x, y) \text{ is } \preceq\text{-minimal in } unprocessed\};$
 $\quad unprocessed \leftarrow unprocessed \setminus stratum;$
 $\quad to_add \leftarrow \emptyset;$
for $case \in stratum$ **do**
 $\quad (case_charac, case_outcome) \leftarrow case;$
 $\quad \text{if the outcome for } case_charac \text{ w.r.t. } (Args, \rightsquigarrow)$
 $\quad \text{is not } case_outcome \text{ then}$
 $\quad \quad to_add \leftarrow to_add \cup \{case\};$
for $case \in to_add$ **do**
 $\quad (Args, \rightsquigarrow) \leftarrow simple_add((Args, \rightsquigarrow), case);$
 $D' \leftarrow Args \setminus \{(\delta_C, \delta_o)\};$
return $D', (Args, \rightsquigarrow)$

case, we test whether it is includable or not w.r.t. the dataset underpinning the current AA framework: if it is, then it is added; otherwise, it is not. More precisely, the algorithm works with strata over D , alongside \preceq . In the simplest setting where each stratum is a singleton, the algorithm works as follows: starting with $D_0 = \emptyset$ and the entire dataset $D = \{d_i\}_{i \in \{1, \dots, |D|\}}$ unprocessed, at each step i , we obtain either $D_i = D_{i-1} \cup \{d_i\}$, if d_i is includable w.r.t. D_{i-1} , and $D_i = D_{i-1}$, otherwise. Then $\hat{D} = D_{|D|} \subseteq D$ is an output of the algorithm. Each example of the current stratum is tested for “includability” with respect to the same (current) subset D_i , and only the includable examples are added to it. Here, however, testing for surprise is enough for this verification. We illustrate the application of the algorithm next.

Example 14. Once more consider the dataset $D = \{(\{a\}, +), (\{c\}, +), (\{a, b\}, +), (\{c, z\}, +), (\{a, b, c\}, +)\}$ in Figure 2c, as well as the definitions used in the proof of Theorem 9 for $X, Y, (\delta_C, \delta_o)$ and \preceq . Let us examine the application of Algorithm 2 to it. We start with an AA framework AF_0 consisting only of

(δ_C, δ_o) , that is, $D_0 = \emptyset$, $AF_0 = AF_{\succeq}(D_0) = AF_{\succeq}(\emptyset) = (\{\{\emptyset, -\}\}, \emptyset)$. The first stratum would be $stratum_1 = (\{\{a\}, +\}, \{\{c\}, +\})$. Of course, then, we have $AA-CBR_{\succeq}(\{\{\emptyset, -\}\}, \{a\}) = -$, and similarly for $(\{c\}, ?)$. Thus, every argument in $stratum_1$ is includable, and is then included in the next AF, resulting in $D_1 = (\{a\}, +), (\{c\}, +)$ and $AF_1 = AF_{\succeq}(D_1)$. Now, the second stratum is $stratum_2 = (\{\{a, b\}, -\}, \{\{c, z\}, -\})$. We can verify that $AA-CBR_{\succeq}(D_1, \{a, b\}) = +$ and $AA-CBR_{\succeq}(D_1, \{c, z\}) = +$. As a result $(\{a, b\}, -)$ and $(\{c, z\}, -)$ are both includable, and then included in the next step, that is, $D_2 = D_1 \cup \{(\{a, b\}, -), (\{c, z\}, -)\}$, and $AF_2 = AF_{\succeq}(D_2)$. Finally, $stratum_3 = (\{\{a, b, c\}, +\})$. Now we verify that $(\{a, b, c\}, +)$ is *not* includable, because $AA-CBR_{\succeq}(D_2, \{a, b, c\}) = +$. Therefore, it is *not* added to the AA framework, that is, $D_3 = D_2$ and thus $AF_3 = AF_{\succeq}(D_3) = AF_{\succeq}(D_2) = AF_2$. Now $unprocessed = \emptyset$, and the selected subset is D_3 , with corresponding $AF_{\succeq}(D_3) = AF_3$, and we are done. Note that using $cAA-CBR_{\succeq}$ the counterexample in the proof of Theorem 9 would fail, since $(\{a, b, c\}, +)$ would not have been added to the AF.

Note that, if D is coherent, we could have defined the algorithm equivalently by looking at cases one-by-one rather than grouping them in strata. However, using strata still has the advantage of allowing for parallel testing of new cases. If D is incoherent, then using strata is necessary.

A full complexity analysis of the algorithm is outside the scope of this paper. However, note here that the algorithm refrains from building the AA framework from scratch each time a new case is considered. Still regarding Algorithm 1, note that it is easy to compute the set DEF while checking whether the next case is includable or not, thus we could optimise its implementation with the use of caching. Besides, the subset of minimal cases (that is, the stratum) can be extracted efficiently by representing the partial order as a directed acyclic graph and traversing this graph. Finally, the order in which the cases in the same stratum are added does not affect the outcome. Thus, each case in the same stratum can be safely tested for includability in parallel.

Definition of $cAA-CBR_{\succeq}$.

Theorem 15. *Let D' be the dataset returned by Algorithm 2. Then for every $\alpha \in D'$, α is surprising w.r.t. D' . Additionally, if D has a concise subset, D' is its unique concise subset. In particular, there is always a concise subset if D is coherent.*

We cannot generalise the existence result for any D : consider the (incoherent) counterexample when $D = (\{\{a\}, +\}, \{\{a, b\}, -\}, \{\{a, b\}, +\})$, for $AA-CBR_{\succeq}$. None of its subsets is concise. Still, our algorithm returns the subset $(\{\{a\}, +\}, \{\{a, b\}, -\})$, which is coherent and consisting only of surprising examples.

To conclude, we can then define inference in $cAA-CBR_{\succeq}$, the classifier yielded by the strategy described until now:

Definition 16. Let D be a dataset and D' be the subset of D identified by Algorithm 2. Let $cAF_{\succeq}(D, N_C)$ be the AF

mined from D' and $(N_C, ?)$, with default argument (δ_C, δ_o) . Then, $cAA-CBR_{\succeq}(D, N_C)$ stands for the outcome for N_C , given $cAF_{\succeq}(D, N_C)$.

Thus, we directly obtain the inference relation $\vdash_{cAA-CBR_{\succeq}}$. Then, $cAA-CBR_{\succeq}$ amounts to the form of $AA-CBR$ using this inference relation. It is easy to see, in line with the discussion before Theorem 13 and using Theorem 8, that $cAA-CBR_{\succeq}$ satisfies several non-monotonicity properties, as follows:

Theorem 17. $\vdash_{cAA-CBR_{\succeq}}$ is cautiously monotonic and also satisfies cut, cumulativity, and rational monotonicity.

Incoherence. An important additional property of $cAA-CBR_{\succeq}$ is that it naturally accommodates a way to handle incoherences in the dataset. During the execution of Algorithm 2, an incoherent pair of cases would be considered at the same stratum. As every characterisation receives an outcome in a $cAA-CBR_{\succeq}$ framework, and exactly one, then if there is an incoherent pair in the dataset, one of its examples would be includable while the other would not. Therefore, only the includable example becomes an argument in the AA framework. Although an incoherent dataset may not have a concise subset, this approach finds a coherent subset which always chooses among one of the conflicting examples, using includability as the criterion for choice.¹⁰ As an example, consider again Figure 1c. Following Algorithm 2, we see that in the first *while* loop both $(\{hm\}, +)$ and $(\{hm\}, -)$ are in the stratum. Since the default outcome is $-$, $(\{hm\}, +)$ is a surprising case w.r.t. \emptyset and thus is added, while $(\{hm\}, -)$ is not and thus is not added, and the algorithm terminates.

Theorem 18. *The dataset returned by Algorithm 2 is coherent.*

Note that since now a coherent subset is used as basis for the inference, whenever the default case is not in the grounded extension, it will be attacked by a case which is in it.¹¹ Thus we have a “principled” way of dealing with incoherences, in which the includable example is always kept.

Spikes. An inconvenience in $AA-CBR_{\succeq}$ is the presence of cases in the AF which do not reach the default case. While part of the AF, they do not affect whether the default case (δ_C, δ_o) is or not in the grounded extension, and thus the outcome. Formally, these cases can be defined as follows, for $AA-CBR_{\succeq}$ as well as $cAA-CBR_{\succeq}$:

Definition 19. Let $(Args, \rightsquigarrow) = AF_{\succeq}(D, N_C)$ or $(Args, \rightsquigarrow) = cAF_{\succeq}(D, N_C)$, and $\alpha \in D \cap Args$. Then, α is a *spike* iff there is no path in $(Args, \rightsquigarrow)$ from α to (δ_C, δ_o) .

As a simple example, consider the casebase in Figure 2c, and add $(\{b\}, -)$ to it. It would be attacked by $(\{a, b, c\}, +)$,

¹⁰Indeed, from this reasoning one can also see that every concise subset is also coherent.

¹¹In more detail, this is so since the AF would be well-founded, and thus every argument outside the grounded extension would be attacked by it (see Dung 1995).

but it would attack no other argument. Thus, $(\{b\}, -)$ would not reach any other argument and would, then, be a spike.

Spikes are unhelpful, since their presence is entirely superfluous, that is, they can be removed with no change in outcome, for any new case.

Theorem 20. *Let $(Args, \rightsquigarrow) = AF_{\succeq}(D, N_C)$ and $\alpha \in D \cap Args$ be a spike. Then $AA-CBR_{\succeq}(D \setminus \{\alpha\}, N_C) = AA-CBR_{\succeq}(D, N_C)$.*

Thus, a useful step in practice is removing spikes from the AF when visualising or storing (e.g. for caching), since the AF may become significantly leaner (indeed, we do this in the case study in Section 8).

Instead, $cAA-CBR_{\succeq}$ shows no spikes, by construction, given that spikes are not includable, and thus are not added to $cAF_{\succeq}(D, N_C)$.

Theorem 21. *Let $(Args, \rightsquigarrow) = cAF_{\succeq}(D, N_C)$. Then, there are no spikes in $Args$.*

8 Case Study

We now explore, as a case study for our approach, the US Trade Secrets domain, frequently discussed in the AI and Law literature (Rissland and Ashley 1987; Brüninghaus and Ashley 2003; Bench-Capon 2017). This area of law deals with misappropriation of commercially relevant information that, allegedly, should not have been available or used by another party. The stereotypical scenario is of a company, the plaintiff, suing another, the defendant, claiming that such misappropriation happened, resulting in economic loss for the plaintiff. In this setting, each case is represented by *factors* each supporting either plaintiff or defendant, and an outcome, which may be a win for plaintiff (Π) or defendant (Δ). Formally, each such case is of the form (F^{Π}, F^{Δ}, o) where F^{Π} are the factors supporting the plaintiff, F^{Δ} , the defendant, and $o \in \{\Pi, \Delta\}$ is the case outcome. Example of pro-plaintiff factors are that the information was about a product which was unique, in the sense that only the plaintiff manufactured this product (F_{Π}^{15}), and that the defendant knew that the information was confidential (F_{Π}^{21}), while some pro-defendant factors are that the plaintiff disclosed the information in negotiations with the defendant (F_{Δ}^1), and that the plaintiff disclosed the information in a public forum (F_{Δ}^{27}). For this case study, we use the publicly available 32 cases considered in (Chorley and Bench-Capon 2005; Al-Abdulkarim, Atkinson, and Bench-Capon 2015; Al-Abdulkarim 2017; Grabmair 2016).

Since factors are polarised representations, that is, they indicate a side, we would lose information in treating them simply as features of $AA-CBR_{\succeq}$. It is necessary to incorporate the idea that, if a case is in favour, for instance, of the plaintiff, then removing one of its pro-defendant factors should still decide the same outcome. This is the idea of a case being *constrained*, as typical in the literature of precedential constraint in AI and Law (Horty and Bench-Capon 2012; Horty 2019; Prakken 2020; Prakken 2021). We accommodate this idea by changing the representation of cases. Formally, if a case is (F^{Π}, F^{Δ}, o) , then it yields the following set of $AA-CBR$ cases: $\{(F^{\Pi} \cup Y, \Pi) \mid Y \subseteq F^{\Delta}\}$, if $o = \Pi$; and $\{(F^{\Delta} \cup Y, \Delta) \mid Y \subseteq F^{\Pi}\}$, if $o = \Delta$. That

is, a single case becomes multiple cases w.r.t. $AA-CBR$. Even though this is not a compact representation (indeed, it is exponential), we only aim to show how $AA-CBR$ and cautious monotonicity applies in this domain, not to provide a scalable representation.

In order to give an appropriate comparison of the resulting AF of both $AA-CBR_{\succeq}$ and $cAA-CBR_{\succeq}$, we remove spikes in the $AA-CBR_{\succeq}$ AF. This makes a more appropriate comparison to $cAA-CBR_{\succeq}$. It turns out that, for this casebase, the resulting AF is the same for both $AA-CBR_{\succeq}$ and $cAA-CBR_{\succeq}$, and shown in Figure 3. However, we show that $AA-CBR_{\succeq}$ could be manipulated by its violation of cautious monotonicity, while $cAA-CBR_{\succeq}$ cannot.

Consider the following new cases $N1_C = (\{F_{\Delta}^1, F_{\Pi}^{18}, F_{\Delta}^{19}, F_{\Pi}^{21}, F_{\Delta}^{23}\}, ?)$ and $N2_C = (\{F_{\Delta}^1, F_{\Delta}^{10}, F_{\Pi}^{18}, F_{\Delta}^{19}, F_{\Pi}^{21}, F_{\Delta}^{23}\}, ?)$. We can think in terms of two cases an attorney needs to argue, and would like to have a specific outcome, for instance, pro-plaintiff (Π). For $N1_C$, the predicted $AA-CBR_{\succeq}$ outcome is Π . For $N2_C$, it is Δ . However, when adding $(N1_C, \Pi)$ to the casebase, the $AA-CBR_{\succeq}$ outcome of $N2_C$ then changes to Π . In terms of the domain, a new case, $N1_C$, which brings no different reason (be it distinguishing, change of social values, among others) for change of the case law, indeed changes the system, as proved by the change in $N2_C$. This implies our attorney in consideration, with no innovation in reasons, could achieve a desired outcome in $N2_C$ by simply presenting it after $N1_C$ is judged. Thus, presenting the cases in different orders would necessarily change the results, even if no new element is introduced, such as considerations of value, policy, or change of legislation. This is not the case for $cAA-CBR_{\succeq}$. It is straightforward to check that $(N1_C, \Pi)$ is non-includable, and thus adding it to the casebase would not change the AF mined from this dataset using $cAA-CBR_{\succeq}$.

9 Related Work and Discussion

Cautious monotonicity is typically discussed in non-monotonic reasoning literature (Makinson 1994; Lehmann and Magidor 1992), originally presented by Gabbay (1984) as a reasonable condition for verifying if an allegedly reasoning system is indeed reasoning, that is, a rationality postulate. It is usually presented along cut and cumulativity, which are argued for by Kraus, Lehmann, and Magidor (1990) on computational and semantic bases.

An important element for the occurrence of incoherence in a dataset is the representation of the cases themselves. That is, an insufficiently expressive knowledge representation risks conflating otherwise distinct cases, giving rise to incoherence if they have different outcomes. We should not think this is a matter left entirely to a human user, which would model a dataset by hand. If cases are thought as being originated from previous processes, such as automatic extraction of features by a natural language processing system, as previously done with $AA-CBR$ (Cocarascu et al. 2020), it is expected that representations could fail in this way, and thus treatment of incoherence is indeed necessary.

There is a long literature on CBR models for legal reasoning, starting with Rissland and Ashley (1987), which

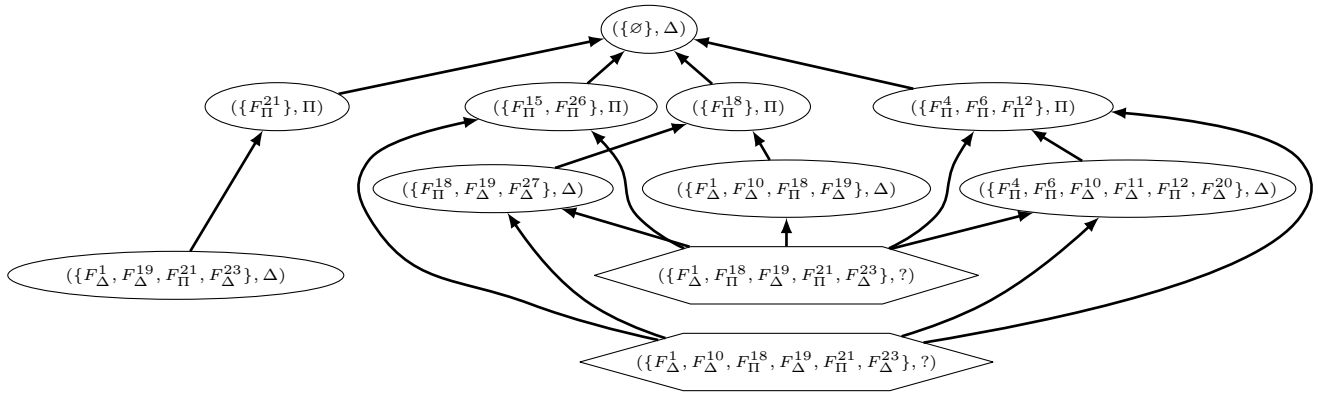


Figure 3: Resulting AA framework for the U.S. Trade Secrets casebase. Each case in the original dataset yields possibly many arguments, each argument represented by its factors and outcomes. Some of the factors are: F_{Δ}^1 : the plaintiff disclosed its product information in negotiations with defendant; F_{Π}^{21} : defendant obtained plaintiff’s information although he knew that plaintiff’s information was confidential F_{Π}^{18} : defendant’s product was identical to plaintiff’s. (See <https://arxiv.org/abs/2107.06413> for a full list of factors.)

is surveyed by Bench-Capon (2017). The original goal of this literature was to capture the argumentative process, and the influence of abstract argumentation on this literature and on AI and Law in general is surveyed by Bench-Capon (2020). However, its goal has expanded to include prediction of cases (Brüninghaus and Ashley 2003; Grabmair 2017) and to explain predictions (Prakken 2020). Our treatment of factors (features for and against) in the case study is non-scalable in general, dictated by the restrictions imposed by the structure of *AA-CBR*. Factors are subject to much research since their appearance in the work of Aleven (2003), and making *AA-CBR* more suitable for dealing with them is a topic left for future work, with argumentation-based treatment of them for CBR already occurring in recent research, such as in the work of Prakken (2020). Another knowledge engineering element also beyond the scope of this work is background knowledge not included in the cases themselves, frequent in the legal CBR literature in the form of a (typically hand-built) domain model, enriching the factor representation (Aleven 2003; Ashley and Brüninghaus 2009; Al-Abdulkarim, Atkinson, and Bench-Capon 2016; Grabmair 2017). For (regular) *AA-CBR*, it is assumed that every relevant knowledge engineering aspect is captured by the partial order, case representations, and default argument.

Notwithstanding this literature, the implications of cautious monotonicity (or the lack of it) to legal reasoning has remained largely unexplored, particularly on CBR scenarios.¹² We illustrate in Section 8 an unexpected consequence of violating it, namely, manipulability of outcomes by leveraging on the order of presentation of new cases. Of course, our analysis is limited and further exploration of the relations between CBR in law and properties of non-monotonic reasoning systems is still required and left to future work.

Horty and Bench-Capon (2012) and Horty (2019) present formal analyses of precedential constraint. In discussing

¹²Prakken (1997) mentions it briefly, and cumulativity is critically discussed in non-monotonic reasoning more generally, but not on a CBR or legally motivated context.

casebase dynamics in the reason model of precedential constraint, Horty and Bench-Capon (2012) found out that “simply following a precedent rule can lead to a change in the law”. One may be led to believe this is an affirmation that an adequate modelling of case law is not cautious monotonic. However, this is not necessarily so. They show that following a rule originated from previous cases may make a decision maker unable to distinguish a new case. That is, merely following a past rule in a case may strengthen the precedential constraint of it, but - and this is the crucial point - we can verify that it would not make a new case previously constrained to an outcome to be constrained to a different outcome. Besides, this effect is only possible if the decision maker is not constrained to an outcome in the changing case (that is, it is still possible to distinguish consistently).

10 Conclusion

We have studied regular *AA-CBR*_≥ frameworks, and proposed a new form of *AA-CBR*, denoted *cAA-CBR*_≥, which is cautiously monotonic and, as a by-product, cumulative and rationally monotonic. We also show that it results in a principled way of dealing with incoherence in casebases, something which *AA-CBR*_≥ lacks. Given that *AA-CBR*_≥ admits the original *AA-CBR*_≥ (Čyras, Satoh, and Toni 2016a) as an instance, we have (implicitly) also defined a cautiously monotonic version thereof.

(Some incarnations of) *AA-CBR* have been shown successful empirically in a number of settings (Cocarascu et al. 2020). The formal properties we have considered in this paper do not necessarily imply better empirical results at the tasks in which *AA-CBR* has been applied. We thus leave for future work an empirical comparison between *AA-CBR*_≥ and *cAA-CBR*_≥. Other issues open for future work are comparisons w.r.t. learnability (such as model performance in the presence of noise), as well as a full complexity analysis of the new model. Also, we conjecture that the reduced size of the AF our method generates could possibly have advantages in terms of time and space complexity: we leave investigation of this issue to future work.

Acknowledgements

We are very grateful to Kristijonas Čyras and Ken Satoh for very valuable discussions, as well as to Alexandre A. Almeida, Victor Nascimento and Matheus Müller for reviewing initial drafts of this paper. We are also grateful for comments by email from Kevin Ashley, as well as for useful comments from anonymous reviewers. The first author was supported by Capes (Brazil, Ph.D. Scholarship 88881.174481/2018-01).

References

- Al-Abdulkarim, L.; Atkinson, K.; and Bench-Capon, T. 2015. Evaluating the use of abstract dialectical frameworks to represent case law. *Proceedings of the 15th International Conference on Artificial Intelligence and Law*.
- Al-Abdulkarim, L.; Atkinson, K.; and Bench-Capon, T. 2016. A methodology for designing systems to reason with legal cases using abstract dialectical frameworks. *Artificial Intelligence and Law* 24(1):1–49.
- Al-Abdulkarim, L. 2017. *Representation of case law for argumentative reasoning*. Ph.D. Dissertation, University of Liverpool, UK.
- Aleven, V. 2003. Using background knowledge in case-based legal reasoning: A computational model and an intelligent learning environment. *Artif. Intell.* 150(1-2):183–237.
- Apt, K. R., ed. 1985. *Logics and Models of Concurrent Systems - Conference proceedings, Colle-sur-Loup (near Nice), France, 8-19 October 1984*, volume 13 of *NATO ASI Series*. Springer.
- Ashley, K. D., and Brüninghaus, S. 2009. Automatically classifying case texts and predicting outcomes. *Artif. Intell. Law* 17(2):125–165.
- Baroni, P.; Cerutti, F.; Giacomin, M.; and Simari, G. R., eds. 2010. *Computational Models of Argument: Proceedings of COMMA 2010, Desenzano del Garda, Italy, September 8-10, 2010*, volume 216 of *Frontiers in Artificial Intelligence and Applications*. IOS Press.
- Bench-Capon, T. J. M. 2017. Hypo’s legacy: introduction to the virtual special issue. *Artif. Intell. Law* 25(2):205–250.
- Bench-Capon, T. J. M. 2020. Before and after Dung: Argumentation in AI and law. *Argument Comput.* 11(1-2):221–238.
- Black, E.; Modgil, S.; and Oren, N., eds. 2015. *Theory and Applications of Formal Argumentation - Third International Workshop, TFAA 2015, Buenos Aires, Argentina, July 25-26, 2015, Revised Selected Papers*, volume 9524 of *Lecture Notes in Computer Science*. Springer.
- Brüninghaus, S., and Ashley, K. D. 2003. Predicting outcomes of case-based legal arguments. In Zeleznikow and Sartor (2003), 233–242.
- Chorley, A., and Bench-Capon, T. J. M. 2005. AGATHA: using heuristic search to automate the construction of case law theories. *Artif. Intell. Law* 13(1):9–51.
- Cocarascu, O.; Stylianou, A.; Čyras, K.; and Toni, F. 2020. Data-empowered argumentation for dialectically explainable predictions. In *ECAI 2020 - 24th European Conference on Artificial Intelligence, Santiago de Compostela, Spain, 10-12 June 2020*.
- Cocarascu, O.; Čyras, K.; and Toni, F. 2018. Explanatory predictions with artificial neural networks and argumentation. In *2nd Workshop on XAI at the 27th IJCAI and the 23rd ECAI*.
1987. *Proceedings of the First International Conference on Artificial Intelligence and Law, ICAIL ’87, Boston, MA, USA, May 27-29, 1987*, ACM.
- Dejl, A.; He, P.; Mangal, P.; Mohsin, H.; Surdu, B.; Voinea, E.; Albini, E.; Lertvittayakumjorn, P.; Rago, A.; and Toni, F. 2021. Argflow: A toolkit for deep argumentative explanations for neural networks. In Dignum, F.; Lomuscio, A.; Endriss, U.; and Nowé, A., eds., *AAMAS ’21: 20th International Conference on Autonomous Agents and Multiagent Systems, Virtual Event, United Kingdom, May 3-7, 2021*, 1761–1763. ACM.
- Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77(2):321 – 357.
- Dung, P. M. 2014. An axiomatic analysis of structured argumentation for prioritized default reasoning. In Schaub et al. (2014), 267–272.
- Dung, P. M. 2016. An axiomatic analysis of structured argumentation with priorities. *Artificial Intelligence* 231:107–150.
- Gabbay, D. M.; Hogger, C. J.; and Robinson, J. A., eds. 1994. *Handbook of Logic in Artificial Intelligence and Logic Programming - Volume 3 - Nonmonotonic Reasoning and Uncertain Reasoning*. Oxford University Press.
- Gabbay, D. M. 1984. Theoretical foundations for non-monotonic reasoning in expert systems. In Apt (1985), 439–457.
- Grabmair, M. 2016. *Modeling purposive legal argumentation and case outcome prediction using argument schemes in the value judgment formalism*. Ph.D. Dissertation, University of Pittsburgh, USA.
- Grabmair, M. 2017. Predicting trade secret case outcomes using argument schemes and learned quantitative value effect tradeoffs. *Proceedings of the 16th edition of the International Conference on Artificial Intelligence and Law*.
- Horty, J. F., and Bench-Capon, T. J. M. 2012. A factor-based definition of precedential constraint. *Artificial Intelligence and Law* 20(2):181–214.
- Horty, J. 2019. Reasoning with dimensions and magnitudes. *Artificial Intelligence and Law* 27(3):309–345.
- Hunter, A. 2010. Base logics in argumentation. In Baroni et al. (2010), 275–286.
- Kenny, E. M., and Keane, M. T. 2019. Twin-systems to explain artificial neural networks using case-based reasoning: Comparative tests of feature-weighting methods in ANN-CBR twins for XAI. In Kraus, S., ed., *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, 2708–2715. ijcai.org.

- Kraus, S.; Lehmann, D.; and Magidor, M. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *Artif. Intell.* 44(1-2):167–207.
- Lehmann, D., and Magidor, M. 1992. What does a conditional knowledge base entail? *Artif. Intell.* 55(1):1–60.
- Makinson, D. 1994. General patterns in nonmonotonic reasoning. In Gabbay et al. (1994), 35–110.
- Nugent, C., and Cunningham, P. 2005. A case-based explanation system for black-box systems. *Artif. Intell. Rev.* 24(2):163–178.
- Paulino-Passos, G., and Toni, F. 2020. Cautious monotonicity in case-based reasoning with abstract argumentation. 18th International Workshop On Non-Monotonic Reasoning.
- Prakken, H.; Wyner, A. Z.; Bench-Capon, T. J. M.; and Atkinson, K. 2015. A formalization of argumentation schemes for legal case-based reasoning in ASPIC+. *J. Log. Comput.* 25(5):1141–1166.
- Prakken, H. 1997. Logical tools for modelling legal argument. *Law and Philosophy Library*.
- Prakken, H. 2020. A top-level model of case-based argumentation for explanation. Proceedings of the ECAI 2020 Workshop on Dialogue, Explanation and Argumentation for Human-Agent Interaction (DEXA HAI 2020).
- Prakken, H. 2021. A formal analysis of some factor- and precedent-based accounts of precedential constraint. *Artificial Intelligence and Law*.
- Rago, A.; Cocarascu, O.; Bechlivanidis, C.; Lagnado, D. A.; and Toni, F. 2021. Argumentative explanations for interactive recommendations. *Artif. Intell.* 296:103506.
- Rissland, E. L., and Ashley, K. D. 1987. A case-based system for trade secrets law. In *Proceedings of the First International Conference on Artificial Intelligence and Law, ICAIL '87, Boston, MA, USA, May 27-29, 1987* (1987), 60–66.
- Schaub, T.; Friedrich, G.; and O’Sullivan, B., eds. 2014. *ECAI 2014 - 21st European Conference on Artificial Intelligence, 18-22 August 2014, Prague, Czech Republic - Including Prestigious Applications of Intelligent Systems (PAIS 2014)*, volume 263 of *Frontiers in Artificial Intelligence and Applications*. IOS Press.
- Shih, A.; Choi, A.; and Darwiche, A. 2019. Compiling bayesian network classifiers into decision graphs. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, 7966–7974.
- Zeleznikow, J., and Sartor, G., eds. 2003. *Proceedings of the 9th International Conference on Artificial Intelligence and Law, ICAIL 2003, Edinburgh, Scotland, UK, June 24-28, 2003*. ACM.
- Čyras, K., and Toni, F. 2015. Non-monotonic inference properties for assumption-based argumentation. In Black et al. (2015), 92–111.
- Čyras, K., and Toni, F. 2016. Properties of ABA+ for non-monotonic reasoning. *CoRR* abs/1603.08714.
- Čyras, K.; Birch, D.; Guo, Y.; Toni, F.; Dulay, R.; Turvey, S.; Greenberg, D.; and Hapuarachchi, T. 2019. Explanations by arbitrated argumentative dispute. *Expert Syst. Appl.* 127:141–156.
- Čyras, K.; Satoh, K.; and Toni, F. 2016a. Abstract argumentation for case-based reasoning. In *KR 2016*, 549–552.
- Čyras, K.; Satoh, K.; and Toni, F. 2016b. Explanation for case-based reasoning via abstract argumentation. In *Proceedings of COMMA 2016*, 243–254.