

Did Karadžić possess the mens rea for genocide in Srebrenica?

Federico Cerutti^{1,2}, Yvonne McDermott Rees³

¹University of Brescia, Department of Information Engineering, via Branze 38, 25123 Brescia, Italy

²Cardiff University, School of Computer Science & Informatics, 5 The Parade, CF24 3AA, Cardiff, UK

³Hillary Rodham Clinton School of Law, Richard Price Building, Swansea University, Singleton Park, Swansea, SA2 8PP, UK

Abstract

We present the methodology and the results of an application of argumentation theory to determine whether Radovan Karadžić, President of the Serb Republic, possessed the *mens rea*—the knowledge of wrongdoing that constitutes part of a crime—for genocide in Srebrenica, where, in July 1995, at least 5,115 Bosnian Muslims were killed by members of the Serb Republic Forces. To evaluate the strengths and weaknesses of Trial Chamber’s findings in the publicly available judgement, we used argumentation-based techniques available in the CISpaces.org tool. The results of our analysis were submitted to the Mechanism for International Criminal Tribunals (MICT) as an *amicus curiae* brief, i.e., a brief from a non-party in a lawsuit who argues or presents information relevant to the lawsuit.

Keywords

argumentation, international criminal law, Srebrenica, genocide

1. Introduction

In this paper, we present the methodology and the results of an application of argumentation theory to determine whether Karadžić possessed *mens rea*¹ for genocide in relation to the Srebrenica mass killing. The results of our analysis were submitted to the Mechanism for International Criminal Tribunals as an *amicus curiae*² brief [1] pursuant to Rule 83 of the MICT Rules of Procedure and Evidence. We based our analysis only on the judgement of Prosecutor v. Radovan Karadžić [2].³

On 24th March 2016, Radovan Karadžić was convicted for genocide in Srebrenica by the International Tribunal for the Prosecution of Persons Responsible for Serious Violations of International Humanitarian Law Committed in the Territory of former Yugoslavia since 1991. As

5th Workshop on Advances In Argumentation In Artificial Intelligence (AI3 2021), part of the 20th International Conference of the Italian Association for Artificial Intelligence (AIXIA 2021).


✉ federico.cerutti@unibs.it (F. Cerutti); yvonne.mcdermottrees@swansea.ac.uk (Y. McDermott Rees)

🌐 <https://federico-cerutti.unibs.it> (F. Cerutti); <https://www.swansea.ac.uk/staff/law/mcdermott-rees-y/>

(Y. McDermott Rees)

🆔 0000-0003-0755-0358 (F. Cerutti); 0000-0003-0111-9049 (Y. McDermott Rees)

© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

¹*Mens rea*: the intention or knowledge of wrongdoing that constitutes part of a crime.

²*Amicus curiae*: a non-party in a lawsuit who argues or presents information relevant to the lawsuit.

³In the following, we will heavily rely upon the judgement [2] as the only source of information for our analysis and paper.

reported in [2], at least 5,115 men were killed by members of the Bosnian Serb Forces in July 1995 in Srebrenica (Section 3).

The Trial Chamber’s finding that the accused possessed the *mens rea*—i.e. the intention and knowledge of wrongdoing that constitutes part of a crime—for genocide in relation to the Srebrenica joint criminal enterprise (JCE) was the subject of academic critique at the time of the Trial judgement, e.g., [3]. Jurisprudence in the International Criminal Court distinguishes genocide from other crimes as it embodies a *dolus specialis*.⁴ A person cannot be treated as complicit in genocide unless they were aware of the *dolus specialis* of the principal perpetrator, hence having the *mens rea*.

Using the argumentation-based techniques available in the CISpaces.org tool [4], reviewed in Section 3, we manually analysed a sub-set of the 2615 pages of [2] to highlight the three reasoning lines that are present in the judgement and that lead to the conclusion that Karadžić possessed the *mens rea*. Of those, two of them might merit further discussions, and the last one relies on a single witness.

Our main contribution is to show that the methodology we propose in Section 4 can be used to show the weakness and strengths of a case—cf. Section 6. This can be of use for the plaintiff, the defendant, but also judges and jurors as it helps clarifying which elements are beyond any reasonable doubt, and which ones are not. This is currently a live issue in international criminal law: one of the authors of this paper argues that “each piece of evidence should be evaluated on its own merits, in light of the other evidence on the record, to determine whether a point has been proven beyond reasonable doubt,” [5] as also supported by several judgements. The opposite is often argued, namely that the Trial Chambers should find their decision on the basis of the the accumulation of all the evidence in the case, but without the need to link factual and legal findings to the final decisions.

The submission of our *amicus curiae* triggered reactions from the academic community interested in international criminal justice, practitioners at the United Nations courts of law, and media. We critically analyse our research and comment on its impact and related work in Section 7.

2. Karadžić and Srebrenica

What follows is a short historical summary of the events that lead to the Srebrenica killings as reported in [2]. It is clearly beyond the scope of this paper to assess the historical truth of those events, hence we will refrain from any comparative study.

The Socialist Republic of Bosnia and Herzegovina (SRBiH) was one of the six republics that once constituted the Socialist Federal Republic of Yugoslavia (SFRY): unlike the other republics, it possessed no single majority ethnic grouping. One of its political parties, the Serbian Democratic Party or SDS—lead by Radovan Karadžić, campaigns to establish separate Serbian institutions. Following a plebiscite held on 9 and 10 November 1991, the Serb Republic was proclaimed in 1992.

Among other key personnel within the Serb Republic, Radovan Karadžić served as President and Supreme Commander of the Bosnian Serb Army (VRS). Tomislav Kovač was the Assistant Minister of the Ministry of Intern (MUP), and the acting Ministry from September 1993 until

⁴*Dolus specialis*: the specific intent to cause a specific kind of harm.

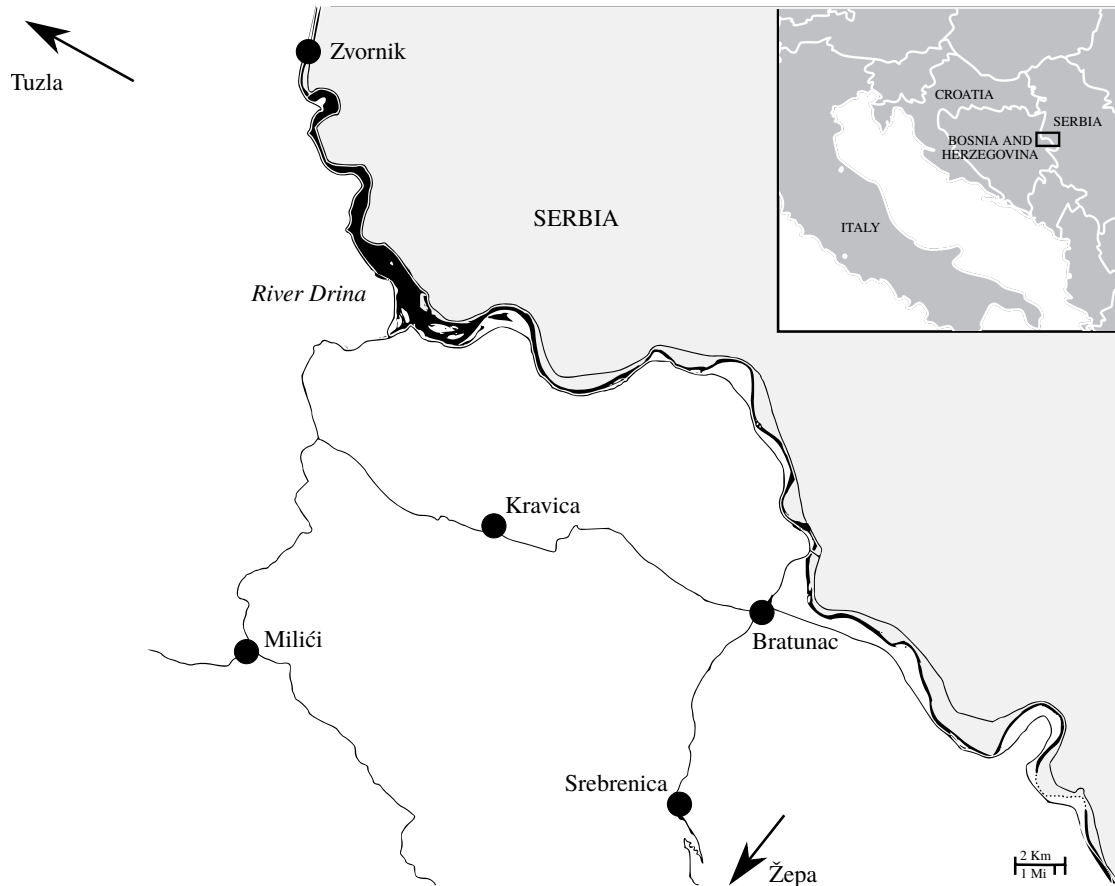


Figure 1: Relevant locations next to the Drina river.

January 1994. Ratko Mladić served as Commander of Main Staff, the highest operative body of the VRS. His assistant commander for Security Administration was Ljubiša Beara, with duties of management of the main staff of the Military Police, as well as co-ordinating with the bodies of the Ministry of the Interior. Momir Nikolić was Chief of the Security and Intelligence Organ, which was responsible for issues of security in the corps composing the VRS, including the arrest and detention of prisoners of war and other persons.

When in 1992 the population of Republic of Bosnia and Herzegovina voted for independence from the Socialist Federal Republic of Yugoslavia in a referendum, forces of the Serb Republic attacked different parts of the Republic of Bosnia and Herzegovina, whose state administration effectively ceased to function having lost control over the entire territory. The Assembly of the Serb Republic adopted the strategic goal to eliminate the border with Serbia: Srebrenica—a town with a majority of Bosnian Muslims—is close to that border (Figure 1).

In late June 1995, Karadžić gave a combat assignment that led to an offensive against Srebrenica and ultimately to the killing of at least 5,115 Bosnian Muslim men.

Figure 2 summarises the timeline of the most relevant events for our analysis starting from the 6th of July.

6th	Shelling of Srebrenica began
11th, Afternoon	Srebrenica has fallen Karadžić appoints Deronjić as Civilian Commissioner for Srebrenica
11th, Night	A column of Bosnian Muslim men started moving in a northwesterly direction towards Tuzla
12th, Morning	Shelling of the column began
12th, Afternoon	Large numbers of the members of the columns surrendered
13th, Morning	Groups of detainees from the column marched towards the Kravica Warehouse
13th, 1630h–Night Kravica Warehouse	One of the Bosnian Muslim detainees took away the rifle of a soldier and shot him dead: other soldiers started shooting at the detainees in response. Others shot at the detainees with machine guns and automatic rifles. Hand-grenades were thrown in the warehouse through the windows. By nightfall, between 755 and 1,016 Bosnian Muslim men were killed.
13th, 1700h–1840h Pale	Karadžić had an hour-long conversation on the phone during which he was briefed by General Maladić that Srebrenica “[wa]s done.”
13th, 2010h Intercepted call between Deronjić and Karadžić through an intermediary	: Deronjić, the President is asking how many thousands? D: About two for the time being. [...] : Deronjić, the President says: “All the goods must be placed inside the warehouses before twelve tomorrow.” D: Right. : Deronjić, not in the warehouses over there, but somewhere else.
13th, around 2100h Bratunac SDS Office	Deronjić ordered to bury the detainees that had been killed at the Kravica Warehouse in a bauxite mine in Milići.
14th Just after midnight	Momir Nikolić drove Beara to the Bratunac SDS office, where Beara met with Deronjić and Vasić. Beara and Deronjić argued about where the Bosnian Muslim men were to be executed, as Beara insisted that he had instructions from his “boss” that the detainees were to remain in Bratunac, and Deronjić countered that the Accused had instructed him that all detainees in Bratunac should be transferred to Zvornik. Eventually, Beara and Deronjić agreed that the detainees would indeed be transferred to Zvornik. Detainees began to be transferred to the first of four detention sites in Zvornik.
14th, 1240h–1310h	Karadžić met with Deronjić alone.
14th, afternoon, after 1310h	Karadžić and Deronjić met with Srebrenica representatives for about four hours.
14th, 2245h–2310h	Kovač met with Karadžić after touring Srebrenica, and the Bratunac and Zvornik areas on 13 and 14 July.
15th, 0035h–0125h	Bajagić—who has a substantive knowledge of the events in Srebrenica being the technical service procurement clerk—met with Karadžić.
16th	By the end of 16th July 1995, at least 3365 Bosnian Muslims men were killed.

Figure 2: Timeline of some of the most relevant events related to the Srebrenica mass killing. All dates refer to the month of July 1995.

3. Background

For this analysis, we used the tools available from the CISpaces project [6] and then further developed in its CISpaces.org version, introduced in [4], that rely on argumentation schemes and computational models of argumentation.⁵

Fundamental concept in computational models of argumentation is the one of inference *rule*, where a statement (*antecedent*) becomes a (*prima facie*) reason to believe another statement (*consequent*). For instance, “Mary, a witness, says that John committed the fraud” (antecedent) can be seen as a *prima facie* reason to believe that “John committed the fraud” (consequent). Some scholars distinguish between strict and defeasible rules in their approach to formal argumentation: as we will not make use of strict rules in this work, we will not discuss this further.

Rules provide the building blocks for the notion of *argument*, that—borrowing from the ASPIC literature [7]—is iterative in the chaining of rules. Statements that are tentatively assumed to hold provide the base case for such an iteration, and thus they are defined as arguments having the statement itself both as premise and as conclusion, where *premises* and *conclusion* are two attributes of the notion of argument. The premises of arguments constructed using this base case also take the name of *ordinary premises* in our approach. Iteratively, an argument requires the existence of a rule whose antecedents are the conclusions of other arguments (*sub-arguments*), and, as a consequent, a statement that becomes the conclusion of this new argument, while its premises are the union of all the premises of its subarguments.

A statement is the contrary of another one when they cannot be both true, albeit they can both be false. Borrowing from the literature, a flexible way of using such a notion of *contrariness* is by allowing for a statement to be the contrary of another one. By requiring the vice versa, the two statements would become *contradictory*. We will make use of such a flexibility in the following of our analysis.

The notion of *contrariness* between statements leads to the concept of *defeat* between arguments: an argument defeats another argument if the former *rebuts* or *undermines* the latter. When the conclusion of an argument contradicts the conclusion of another argument, it is the case that the first *rebuts* the second, as well as all the other arguments that have such a second argument as sub-argument. If, instead, the conclusion of an argument contradicts one of the ordinary premises of another one, then the former *undermines* the latter.

Given a set of arguments and defeats between them, we need criteria to assess which arguments collectively survive the defeats and thus can provide a reasonable viewpoint (or *extension*) of the statements and the rules that we were considering. Such criteria usually consider *conflict-freeness*, i.e. the absence of defeats within the viewpoint; *admissibility*, i.e. if an argument in the viewpoint is defeated by a second argument, the latter must in turn be defeated by a third argument also in the viewpoint; and *maximality*, i.e. a viewpoint cannot be a strict subset of another viewpoint. Multiple viewpoints can exist for the same set of arguments and the defeats between them: two equally reliable witnesses, each providing one reason for contradictory conclusions, lead to the situation that each of the two arguments per se is a reasonable viewpoint, hence there are two of them. In this case, if an argument belongs to at least one viewpoint, it is said to be *credulously accepted*. If, instead, an argument belongs to all the viewpoints, it is said to be *sceptically*

⁵An interested reader can try the CISpaces.org at <https://tiresia.unibs.it/cispaces/>.

accepted. In the following we will be making use of this notion of sceptically acceptance in connection with the principle of *proof beyond a reasonable doubt*.

CISpaces.org provides a convenient visual language and an effective Human-Machine Interface for argumentation mapping. It builds on top of the Argument Interchange Format AIF [8] that specifies a graph structure composed of two types of nodes connected by links. The nodes can be either information (in the following identified by squared boxes) or scheme nodes (in the following identified by round boxes). Information nodes define the antecedents and consequents that we will be making use in the generation of arguments. Scheme nodes can be either rule of inference nodes or conflict nodes. An inference node provides the connection between antecedents and a consequent: if one or more information nodes are linked to an inference node, and the latter is in turn linked to another information node, we will interpret this sub-graph as an inference rule. Conflict nodes, instead, express the contrariness relationship between two inference nodes: once again, links here are directed too.

3.1. Argumentation Schemes

Argumentation schemes [9] are abstract reasoning patterns commonly used in everyday conversational argumentation, legal, scientific argumentation, etc. Schemes have been derived from empirical studies of human arguments and debate. Each scheme has a set of critical questions that represent standard ways of critically probing into an argument to find aspects of it that are open to criticism. For instance, the following is the scheme for arguments from *evidence to hypothesis* [9]:

Major Premise: If A (a hypothesis) is true, then B (a proposition reporting an event) will be observed to be true.

Minor Premise: B has been observed to be true in a given instance.

Conclusion: Therefore, A is true.

CQ1: Is it the case that if A is true, then B is true?

CQ2: Has B been observed to be true?

CQ3: Could there be some reason why B is true, other than its being true because of A being true?

The other argumentation schemes used in this analysis are: the *abductive argumentation scheme*; the *argumentation from cause to effect*; the *argumentation from (popular) opinion*; and the *argumentation from witness testimony* [9].

3.2. Charting and Evaluation of Arguments

CISpaces.org [6, 4] enables a user to draw a directed graph ($WDG = \langle N, \rightsquigarrow \rangle$) based on the AIF format [8], thus with two distinct types of nodes: information nodes (or I-nodes) and scheme nodes (or S-nodes). S-nodes can be either rule of inference application (RA-nodes), or conflict

application (CA-nodes), respectively represented as *Pro* and *Con* nodes. Pro links can be further labelled with the argumentation scheme they instantiate. In a *WDG*, nodes are connected by edges whose semantics are implicitly defined by their use [8].

Similarly to [7, 10], a *WDG* can be mapped into an ASPIC+ system [11]. Assume a logical language \mathcal{L} , and a set of *strict* or *defeasible* inference rules—resp. $\varphi_1, \dots, \varphi_n \longrightarrow \varphi$ and $\varphi_1, \dots, \varphi_n \Longrightarrow \varphi$. A strict rule inference always holds—i.e. if the *antecedents* $\varphi_1, \dots, \varphi_n$ hold, the *consequent* φ holds as well—while a defeasible inference “usually” holds.

An *argumentation system* is as tuple $AS = \langle \mathcal{L}, \mathcal{R}, \bar{\cdot}, \nu \rangle$ where:

- $\bar{\cdot} : \mathcal{L} \mapsto 2^{\mathcal{L}}$ is a contrariness function s.t. if $\varphi \in \bar{\psi}$ and: $\psi \notin \bar{\varphi}$, then φ is a *contrary* of ψ ; $\psi \in \bar{\varphi}$, then φ is a *contradictory* of ψ ($\varphi = -\psi$);
- $\mathcal{R} = \mathcal{R}_d \cup \mathcal{R}_s$ is a set of strict (\mathcal{R}_s) and defeasible (\mathcal{R}_d) inference rules such that $\mathcal{R}_d \cap \mathcal{R}_s = \emptyset$;
- $\nu : \mathcal{R}_d \mapsto \mathcal{L}$, is a partial function.⁶

A *knowledge base* in an AS is a set $\mathcal{K}_n \cup \mathcal{K}_p = \mathcal{K} \subseteq \mathcal{L}$; $\{\mathcal{K}_n, \mathcal{K}_p\}$ is a partition of \mathcal{K} ; \mathcal{K}_n contains *axioms* that cannot be attacked; \mathcal{K}_p contains *ordinary premises* that can be attacked. An *argumentation theory* is a pair $AT = \langle AS, \mathcal{K} \rangle$.

To map a *WDG* into an ASPIC+ system, let us assume that:

- $P \subseteq N$ is the set of I-nodes, where each I-node in the graph is written p_i ;
- ℓ_{type} , with $type = \{Pro, Con\}$, refers to a S-node;
- $[p_1, \dots, p_n \rightsquigarrow \ell_{Pro} \rightsquigarrow p_\phi]$ indicates an inference rule, where p_1, \dots, p_n are parent nodes of the S-node ℓ_{Pro} , and p_ϕ is a child of ℓ_{Pro} ;
- conflict schemes can be either $[p_1 \rightsquigarrow \ell_{Con} \mapsto p_2]$ or $[p_1, \dots, p_n \rightsquigarrow \ell_{Con} \rightsquigarrow p_\phi]$.

Given a *WDG* $\langle N, \rightsquigarrow \rangle$, its corresponding ASPIC+ system $AS = \langle \mathcal{L}, \mathcal{R}, \bar{\cdot}, \nu \rangle$ is such that:

- $\forall p \in P \subseteq N, p \in \mathcal{L}$;
- $\mathcal{R} = \mathcal{R}_s \cup \mathcal{R}_d$ with $\mathcal{R}_s = \emptyset$ and $\forall [p_1, \dots, p_n \rightsquigarrow \ell_{Pro} \rightsquigarrow p_\phi], p_1, \dots, p_n \Rightarrow p_\phi \in \mathcal{R}_d$;
- $\forall [p_1 \rightsquigarrow \ell_{Con} \rightsquigarrow p_2], p_1 \in \bar{p_2}$;
- $\forall [p_1, \dots, p_n \rightsquigarrow \ell_{Con} \rightsquigarrow p_\phi]$, is mapped as $p_1, \dots, p_n \Rightarrow p_h \in \mathcal{R}_d$ and $p_h \in \bar{p_\phi}$;

and the knowledge base $\mathcal{K}_n \cup \mathcal{K}_p = \mathcal{K} \subseteq \mathcal{L}$ is such that, given $[p_1, \dots, p_n \rightsquigarrow \ell_{Pro} \rightsquigarrow p_\phi]$, $\forall p_i \in \{p_1, \dots, p_n\}$, if p_i is not a conclusion of any inference rule $\exists [\ell_{Pro} \rightsquigarrow p_i] \in WDG$, $p_i \in \mathcal{K}_p$. In addition, assume $WDG' = \langle N', \rightsquigarrow' \rangle$ a subset of *WDG* containing only a single cycle of inference schemes, then $\forall p_i \in P' \subset N'$, if $[\ell_{Pro} \rightsquigarrow p_i], [p_i \rightsquigarrow \ell_{Pro}] \in WDG'$, then $p_i \in \mathcal{K}_p$.

Following [11], an *argument* \mathbf{a} on the basis of a $AT = \langle AS, \mathcal{K} \rangle$, $AS = \langle \mathcal{L}, \mathcal{R}, \bar{\cdot}, \nu \rangle$ is:

1. φ if $\varphi \in \mathcal{K}$ with: $\text{Prem}(\mathbf{a}) = \{\varphi\}$; $\text{Conc}(\mathbf{a}) = \varphi$; $\text{Sub}(\mathbf{a}) = \{\varphi\}$; $\text{Rules}(\mathbf{a}) = \text{DefRules}(\mathbf{a}) = \emptyset$; $\text{TopRule}(\mathbf{a}) = \text{undefined}$.

⁶Informally, $\nu(r)$ is a wff in \mathcal{L} which says that the defeasible rule r is applicable. However, we will not make use of this feature in the following.

2. $\mathbf{a}_1, \dots, \mathbf{a}_n \rightarrow / \implies \psi$ if $\mathbf{a}_1, \dots, \mathbf{a}_n$, with $n \geq 0$, are arguments such that there exists a strict/defeasible rule $r = \text{Conc}(\mathbf{a}_1), \dots, \text{Conc}(\mathbf{a}_n) \rightarrow / \implies \psi \in \mathcal{R}_s/\mathcal{R}_d$. $\text{Prem}(\mathbf{a}) = \bigcup_{i=1}^n \text{Prem}(\mathbf{a}_i)$; $\text{Conc}(\mathbf{a}) = \psi$; $\text{Sub}(\mathbf{a}) = \bigcup_{i=1}^n \text{Sub}(\mathbf{a}_i) \cup \{\mathbf{a}\}$; $\text{Rules}(\mathbf{a}) = \bigcup_{i=1}^n \text{Rules}(\mathbf{a}_i) \cup \{r\}$; $\text{DefRules}(\mathbf{a}) = \{d \mid d \in \text{Rules}(\mathbf{a}) \cap \mathcal{R}_d\}$; $\text{TopRule}(\mathbf{a}) = r$

An argument can be attacked in its premises (*undermining*) or its conclusion (*rebuttal*). Since we will not use the preference ordering between arguments, we will omit it from the definition. Similarly for the notion of *undercut* on the inference rule (cf. [11]).

Given \mathbf{a} and \mathbf{b} arguments, \mathbf{a} *defeats* \mathbf{b} iff \mathbf{a} *successfully rebuts* or *successfully undermines* \mathbf{b} , where: \mathbf{a} *successfully rebuts* \mathbf{b} (on \mathbf{b}') iff $\text{Conc}(\mathbf{a}) \notin \bar{\varphi}$ for some $\mathbf{b}' \in \text{Sub}(\mathbf{b})$ of the form $\mathbf{b}'_1, \dots, \mathbf{b}'_n \implies -\varphi$; \mathbf{a} *successfully undermines* \mathbf{b} (on φ) iff $\text{Conc}(\mathbf{a}) \notin \bar{\varphi}$, and $\varphi \in \text{Prem}(\mathbf{b}) \cap \mathcal{K}_p$.

An *argumentation framework* (AF) [12] is a pair $\Delta = \langle \mathcal{A}, \rightarrow \rangle$ where \mathcal{A} is a set of arguments⁷ and $\rightarrow \subseteq \mathcal{A} \times \mathcal{A}$. We denote with $\mathbf{a}_2 \rightarrow \mathbf{a}_1$ when $\langle \mathbf{a}_2, \mathbf{a}_1 \rangle \in \rightarrow$.

An AF $\langle \mathcal{A}, \rightarrow \rangle$ is the *abstract argumentation framework* defined by $AT = \langle AS, \mathcal{K} \rangle$, $AS = \langle \mathcal{L}, \mathcal{R}, \bar{\cdot}, \nu \rangle$ if \mathcal{A} is the smallest set of all finite arguments constructed from \mathcal{K} (as above); and \rightarrow is the defeat relation on \mathcal{A} .

Given an AF $\Delta = \langle \mathcal{A}, \rightarrow \rangle$: a set $S \subseteq \mathcal{A}$ is a *conflict-free* set of Δ if $\nexists \mathbf{a}_1, \mathbf{a}_2 \in S$ s.t. $\mathbf{a}_1 \rightarrow \mathbf{a}_2$; an argument $\mathbf{a}_1 \in \mathcal{A}$ is *acceptable* with respect to a set $S \subseteq \mathcal{A}$ of Δ if $\forall \mathbf{a}_2 \in \mathcal{A}$ s.t. $\mathbf{a}_2 \rightarrow \mathbf{a}_1$, $\exists \mathbf{a}_3 \in S$ s.t. $\mathbf{a}_3 \rightarrow \mathbf{a}_2$; a set $S \subseteq \mathcal{A}$ is an *admissible* set of Δ if S is a conflict-free set of Δ and every element of S is acceptable with respect to S .

A set of argument $S \subseteq \mathcal{A}$ is a *preferred extension* iff it is a maximal (w.r.t. set inclusion) admissible set.

An argument is *skeptically accepted* w.r.t. preferred semantics iff it belongs to *each* preferred extension. Checking this is a problem that lies at the second level of the polynomial hierarchy [14], hence the need—in general—for efficient implementations [15].

4. Methodology

The goal of our *amicus curiae* brief [1] was to identify the precise factual and inferential bases for the Trial Chamber's findings of Karadžić's genocidal intent in the Trial Chamber judgement, and to elucidate the forms of reasoning that led to these conclusions. We limited our analysis to the reasoning process that can be fathomed from the Trial Chamber's judgment. As such, we did not analyse issues such as the reliability of witnesses or evidence: they are the purview of the Trial Chamber alone, also because the entire set of evidence used by the Trial Chamber is not publicly available.

In the present case, Karadžić's *mens rea* is an element of the offence of genocide in Srebrenica, as genocide requires each member of the joint criminal enterprise to be knowledgeable of the *dolus specialis* of the principal perpetrator. The *material facts* upon which proof of *mens rea* hinged were the Trial Chamber judgement's findings on: (1) Karadžić's knowledge of the expansion of the plan to remove Bosnian Muslims from Srebrenica to include the killing of men and boys, hence Karadžić sharing the intent to destroy the Bosnian Muslims in Srebrenica; and (2) his active involvement in the killings. Due to space limitations, let us consider only the former.

⁷In this paper we consider only *finite* sets of arguments: see [13] for a discussion on infinite sets of arguments.

Following [16], we manually and unanimously identified the arguments—and their general argumentation schemes when possible—that the Trial Chamber put forward in [2] related to the given hypothesis, together with (1) those instances of schemes for which not all critical questions have been satisfactorily addressed; (2) and particular facts that seem missing but necessary to expose the entire line of reasoning, labelling them with **Unstated**. In those cases, we did not include an analysis of critical questions for the inferences based on such unstated pieces of information, as it would be a detour from the purpose of this work.

It is worth noticing that the text proved resistant to attempts of automatic analysis. This is also evident in the graphical charting of our analysis (Figure 3), where we consider pieces of information spanning more than 210 pages (Para 5312 fn. 18025 [2, p. 2203] to Para 5808 [2, p. 2413]), in addition to historical information scattered around the entire document. For instance, the information that Srebrenica has fallen on 11 July 1995 has been presented in Para 5033 [2, p. 2079], 331 pages before being used in an argument to support the hypothesis.

5. Results

Figure 3 depicts our understanding of the reasoning lines that the Trial Chamber describes in its judgement [2] in support of the hypothesis that Karadžić was knowledgeable of the intent of killing Bosnian Muslims men. This is also the conclusion of the skeptically accepted arguments w.r.t. preferred semantics (cf. Section 3), quite unsurprisingly given the scarce number of conflicts: this is expected since the judgement does not record each exchange of arguments between the defence and the prosecution.

There are three main lines of reasoning in favour of this conclusion. The first one is based upon Nikolić's testimony that he overheard Deronjić saying that the accused had instructed Deronjić that all detainees should be transferred to Zvornik, cf. Figure 2, 14th July 1995, just after midnight. This testimony gives reasons to the chamber to refute the alternative explanation—highlighted by the defence and reported in the judgement—that Karadžić was referring to a place different from Zvornik in the intercepted conversation with Deronjić, cf. Figure 2, 13th July 1995, 2010h. In this line of reasoning, the chamber decided also to link additional pieces of information (*Inference 3.cX* of Figure 3), as supporting the conclusion that Karadžić ordered that detainees should be transferred to Zvornik, such as a complaint to Beara by Deronjić about the presence of detainees in Bratunac. However, for those facts, it seems to us that the chamber did not consider some relevant critical question, e.g. *Is there any other reasonable explanation for why Deronjić had previously complained to Beara about the detainees' presence in Bratunac, other than it being true because Karadžić conveyed to Deronjić the direction that the detainees should be transferred to Zvornik?* Despite those additional pieces of information (*Inference 3.cX* of Figure 3), this line of reasoning does not rely on unstated findings or pieces of information for which critical questions have not explicitly been answered. Let us remind that we methodologically chose not to assess the reliability of Nikolić testimony as we did not have access to the entire historical records.

A second line of reasoning justifying the hypothesis is based on Simić's testimony that Deronjić told him that he had informed Karadžić about the events at the Kravica Warehouse the day after the incident, in conjunction with the unstated assumption that if the accused knew that Bosnian

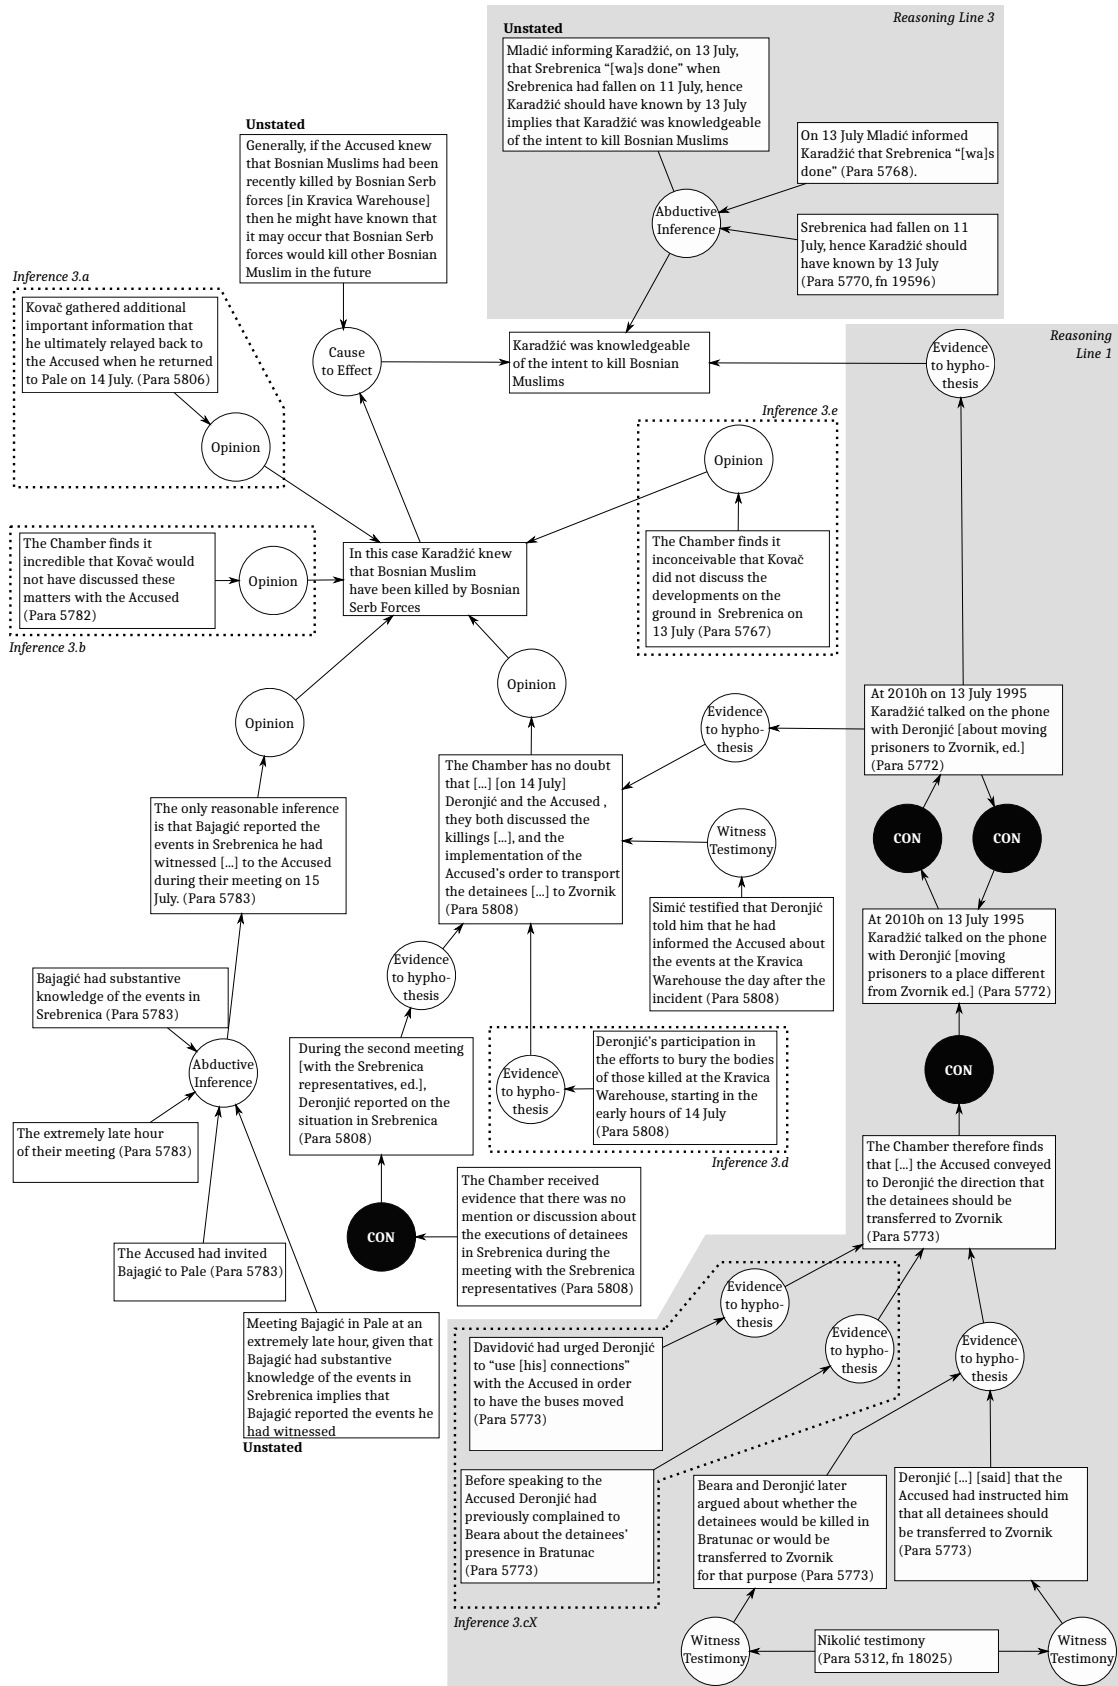


Figure 3: Analysis of arguments in [2] in favour of the hypothesis that Karadžić was knowledgeable of the intent of killing Bosnian Muslims men. Each Para reference refers to [2]. Names and events are introduced in Section 2, except for Milorad Davidović, who was a senior official in the MUP and later on a witness. Squared boxes are claims; white circles inferences (Pro), eventually labelled with the argumentation schemes they refer to; while black circles conflict relations (Con). Dotted areas identify inferences for which there are critical questions that were not explicitly addressed in [2]. Three reasoning lines are highlighted as they are referred to in Section 5.

Muslims had been recently killed by Bosnian Serb forces (in Kravica Warehouse), then he might have been known that it may occur that Bosnian Serb forces would kill other Bosnian Muslims in the future. Similarly as above, we have some doubts that all relevant critical questions find an answer in the judgement [2], w.r.t. *Inferences 3.a, 3.b, 3.d, 3.e* of Figure 3. For instance, *what evidence supported the finding that Kovač relayed back additional important information to Karadžić when he returned to Pale on 14 July? (Inference 3.a, Figure 3); or what evidence supported the finding that Kovač discussed these matters with Karadžić? (Inference 3.b, Figure 3).*

A third line of reasoning is based on an abductive inference with the unstated premise that Mladić informed Karadžić, on 13 July, that Srebrenica “[wa]s done.” The Trial Chamber appears to have concluded that, given that Srebrenica had fallen on 11 July, Karadžić would have known this by 13 July. From that unstated inference, it drew a further inference that the conversation implied that Karadžić knew of the intent to kill the Bosnian Muslims of Srebrenica.

6. Discussion

On 21 February 2018, we sent to the United Nations Mechanism for International Criminal Tribunals (MICT) a request for leave to make submissions as *amicus curiae* pursuant to Rule 83 of the MICT Rules of Procedure and Evidence [1].

Given the overall results (cf. Section 5), we believed that our analysis was particularly useful for the prosecutor in the case, as it highlights a few ways to strengthen the case. Much to our surprise, on 5 March 2018 the prosecution replied to our request asking the Appeals Chamber to deny the motion because *amicus curiae* submissions are *generally* limited to questions of law, and not to questions of facts. On 9 March we then replied that whilst our *amicus curiae* brief focuses on the question of Karadžić’s *mens rea* for genocide in Srebrenica, it goes to the heart of an important legal question on the interpretation of the standard of proof, i.e., whether each piece of evidence should be evaluated on its own merit or whether it is acceptable for the Trial Chambers to find their decision on the basis of the accumulation of all the evidence in the case but without the need to link factual and legal findings to the final decision.

With an even greater surprise, on 12 March, the Defence Counsel asked the Appeals Chamber to accept the proposed submissions and give the parties the opportunity to make their own observations on those submissions, adding that “[t]he Proposed Submissions bring judicial fact-finding into the 21st century.”

This is clearly not the first attempt to bring “judicial fact-finding into the 21st century.” In [17], Verheij introduces the notion of *automated argument assistance* which is in spirit very close to our work here, as it explicitly aims at drafting and testing of court pleadings. Walton [18] provides an extensive account of argumentation in legal systems, and argues for a *new method* for legal argumentation which, among other, includes the use of “argument diagramming to map out the network of inference in a given case” [18, p. 323]. For completeness of discussion, authoritative colleagues criticise the use of argument diagramming, notably van Gelder in [19]. In reflecting on his experience, he notices that argument diagramming might not serve well the purpose of deliberation, possibly because deliberation is a dialectical activity rich in nuances. However, he did not consider deliberation activities where the incentives for a proper epistemic investigation

are significant, such as writing a judgement for an international criminal case.

The work by Walton on legal argumentation [18] and in general on argumentation schemes—summarised in [9]—motivated researchers in deriving computational models, thus building on the tradition initiated by Verheij [17]. Bex et.al. [20] expanded on the idea of using argumentation schemes for providing a formal account of reasoning, and subsequently in [21] they also considered the advantages of merging it with storytelling. The latter also takes into consideration the different positions of the plaintiff and the defendant, which is also the case of [22]—where a formal dialogue system is used as a formalisation tool—and [23], where ASPIC+ is used for formalising legal case-based reasoning. Differently from the previous approaches, we considered explicitly the role of argumentative semantics using skeptical preferred as a proxy for the *beyond any reasonable doubt* standard of proof. This is clearly questionable, but it looks a reasonable approximation as it is a rather conservative choice, although it might be a little difficult to explain to non-experts. Further analysis using other semantics are already envisaged, as well as a deeper comparison with the ANGELIC methodology [24], in particular after the recent paper [25] showing a correspondence with ASPIC+.

We also feel that there is very little we can add to van Gelder observations in [26], where he analyses some legal arguments. His comments strongly resonate with us, as we also experienced “little use of verbal indicators of logical structure, and often use obscure or vague indicators,” [26] incomplete arguments, with also text scattered across the document, and possibly serving multiple purposes.⁸ Although far from providing an off-the-shelf support tool, in retrospective our analysis would have significantly benefit from (1) entity-relations extractors, e.g. [27] and (2) topic modelling system, e.g. [28], which together might transform a static PDF document into a database that can be queried. As it is evident in the graphical charting of our analysis (Figure 3), we considered pieces of information scattered across more than 210 pages (Para 5312 fn. 18025 [2, p. 2203] to Para 5808 [2, p. 2413]), in addition to historical information written in different part of the document: for instance, the information that Srebrenica has fallen on 11 July 1995 has been presented in Para 5033 [2, p. 2079], 331 pages before being used in an argument.

Differently from previous approaches, we considered a case under discussion at ICTY offering the results of our analysis as an *amicus curiae* brief to the Appeal Chambers. It unfortunately denied admissibility of our application on 28 March 2018, observing that “the issues regarding whether Karadžić the *mens rea* for genocide in relation to the Srebrenica JCE were extensively litigated before the Trial Chamber and have been fully briefed by Karadžić and the Prosecution on appeal.” The Appeals Chamber also seems to criticise the fact that “the *Amicus Curiae* Observations seek to guide the Appeals Chamber’s analysis of the Trial Judgement without consideration of or access to the entire record that is relevant to the Trial Chamber’s conclusions.” This however would raise the question: what is the purpose of having a 2615 page judgement, if the judgment does not actually fully reflect the grounds for the conclusion? Finally, confirmation that our analysis was trustworthy comes from the Appeal Chamber Judgement [29] that in its section D.2 provides a summary of the Trial’s Chamber Judgement regarding whether Karadžić was knowledgeable of the intent to kill Bosnian Muslims which is almost entirely present in our

⁸A reviewer of an earlier version of this paper commented that some of the instance of argument from opinion in Figure 3 seem more of instance of argument from ignorance. This is something that only a judge mindful of the purpose of their prose could clarify when writing the document.

resulting argumentation network depicted in Figure 3.

7. Conclusion

In this paper, we present the methodology and the results of an application of argumentation theory to determine whether Karadžić possessed *mens rea* for genocide in Srebrenica based on [2]. As discussed in Section 5, we summarise the results of our analysis testing whether Karadžić was knowledgeable of the intent—of General Mladić and others—to kill Bosnian Muslims. This hypothesis is supported within the Trial Chamber’s judgement [2] by three lines of reasoning, two of which might merit further discussion, and the last one relying on a single witness.

Although at first sight this paper seems to be similar to other attempts to analyse legal reasoning with formal argumentation, e.g. [22, 23], it differs from them substantially as we did not try to capture the debate component, hence distinguishing between Prosecutor and Defence claims. Instead, our analysis is closer to works analysing arguments in a single document, like, for instance, [16] that analyses the role of argumentation in written financial communications.

Although the Appeals Chamber denied the admissibility of our application, the interest that applying formal argumentation theories triggered in the international criminal law community suggests that there is scope for future work in this area. We cannot claim that the methodology used in this analysis is beyond critique, but we can claim that it can help creating a better judgement that fully reflects the grounds for the overall conclusion.

This is the long-term aspiration of the ongoing research underpinning this paper, and we are fully aware that this will require to provide answers and innovative proposals both from a technical perspective as well as from the legal one. From a technical perspective, for instance, we still lack appropriate methodologies for adequately transforming statements of natural language into formal logic—a problem most students of logic encounter without being presented with satisfactory solutions, cf. among others [30]—thus inevitably exposing the subjectivity of each formalisation. In addition, following [31], we will also work in the direction of assessing the quality and the strengths of different argumentation reasoning lines, by taking into consideration quantitative measurements of uncertainty and trust, thus enriching the community proposals looking at probabilistic elements in legal reasoning, e.g. [32, 33, 34].

Acknowledgments

This research formed part of a project entitled Developing a Tool to Support Judicial Reasoning in International Criminal Trials, funded in part by Cherish-DE, the United Kingdom’s Digital Economy Crucible (<http://cherish-de.uk>).

This research was partially sponsored by the Italian Ministry of Research through a Rita Levi-Montalcini Personal Fellowship (D.M. n. 285, 29/03/2019) and by the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-16-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K.

Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

- [1] Y. McDermott Rees, F. Cerutti, Request for leave to make submissions as amicus curiae, <http://jrad.unmict.org/webdrawer/webdrawer.dll/webdrawer/rec/240941/view/>, 2018.
- [2] ICTY, Prosecutor v. Karadžić. Judgement, http://www.icty.org/x/cases/karadzic/tjug/en/160324_judgement.pdf, 2016.
- [3] M. Sterio, The Karadžić Genocide Conviction: Inferences, Intent, and the Necessity to Redefine Genocide, *Emory International Law Review* (2017).
- [4] F. Cerutti, T. J. Norman, A. Toniolo, S. E. Middleton, Cispaces.org: from fact extraction to report generation, in: *COMMA 2018*, 2018, pp. 269 – 280.
- [5] Y. McDermott, Inferential reasoning and proof in international criminal trials: The potentials of wigmorean analysis, *Journal of International Criminal Justice* 13 (2015) 507–533.
- [6] A. Toniolo, T. J. Norman, A. Etuk, F. Cerutti, R. W. Ouyang, M. Srivastava, N. Oren, T. Dropps, J. A. Allen, P. Sullivan, Agent Support to Reasoning with Different Types of Evidence in Intelligence Analysis, in: *AAMAS 2015*, 2015, pp. 781–789.
- [7] H. Prakken, An abstract framework for argumentation with structured arguments, *Argument & Computation* 1 (2010) 93–124.
- [8] C. I. Chesnevar, J. McGinnis, S. Modgil, I. Rahwan, C. Reed, G. R. Simari, M. South, G. A. W. Vreeswijk, S. Willmot, Towards an argument interchange format, *The Knowledge Engineering Review* 21 (2006) 293.
- [9] D. Walton, C. Reed, F. Macagno, *Argumentation schemes*, Cambridge University Press, NY, 2008.
- [10] M. Janier, J. Lawrence, C. Reed, OVA+: An argument analysis interface, in: *COMMA 2014*, 2014, pp. 463–564.
- [11] S. Modgil, H. Prakken, A general account of argumentation with preferences, *Artificial Intelligence* 195 (2013) 361–397.
- [12] P. M. Dung, On the Acceptability of Arguments and Its Fundamental Role in Nonmonotonic Reasoning, Logic Programming, and n-Person Games, *Artificial Intelligence* 77 (1995) 321–357.
- [13] P. Baroni, F. Cerutti, P. E. Dunne, M. Giacomin, Automata for Infinite Argumentation Structures, *Artificial Intelligence* 203 (2013) 104–150.
- [14] P. E. Dunne, M. Wooldridge, Complexity of abstract argumentation, in: *Argumentation in AI*, Springer-Verlag, 2009, pp. 85–104.
- [15] F. Cerutti, M. Vallati, M. Giacomin, An Efficient Java-Based Solver for Abstract Argumentation Frameworks: jArgSemSAT, *International Journal on Artificial Intelligence Tools* 26 (2017) 1750002.
- [16] R. Palmieri, A. Rocci, N. Kudrautsava, Argumentation in earnings conference calls. corporate standpoints and analysts’ challenges, *Studies in Communication Sciences* 15 (2015) 120 – 132.

- [17] B. Verheij, Bart, Automated argument assistance for lawyers, in: ICAIL '99, New York, New York, USA, 1999, pp. 43–52.
- [18] D. N. Walton, Legal argumentation and evidence, Pennsylvania State University Press, 2002.
- [19] T. van Gelder, Cultivating Deliberation for Democracy, *Journal of Public Deliberation* 8 (2012).
- [20] F. Bex, H. Prakken, C. Reed, D. Walton, Towards a Formal Account of Reasoning about Evidence: Argumentation Schemes and Generalisations, *Artificial Intelligence and Law* 11 (2003) 125–165.
- [21] F. Bex, S. van den Braak, H. van Oostendorp, H. Prakken, B. Verheij, G. Vreeswijk, Sense-making software for crime investigation: how to combine stories and arguments?, *Law, Probability and Risk* 6 (2007) 145–168.
- [22] H. Prakken, Formalising ordinary legal disputes: a case study, *Artificial Intelligence and Law* 16 (2008) 333–359.
- [23] H. Prakken, A. Wyner, T. Bench-Capon, K. Atkinson, A formalization of argumentation schemes for legal case-based reasoning in ASPIC+, *Journal of Logic and Computation* 25 (2015) 1141–1166.
- [24] L. Al-Abdulkarim, K. Atkinson, T. J. M. Bench-Capon, A methodology for designing systems to reason with legal cases using abstract dialectical frameworks, *Artificial Intelligence and Law* 24 (2016) 1–49.
- [25] K. Atkinson, T. J. M. Bench-Capon, Relating the ANGELIC methodology and ASPIC+, in: *COMMA 2018*, 2018, pp. 109–116.
- [26] T. van Gelder, Why are legal arguments so hard to follow?, <https://timvangelder.com/2009/08/13/why-are-legal-arguments-so-hard-to-follow/> (on 23 Aug 2018), 2009.
- [27] M. Surdeanu, D. McClosky, M. R. Smith, A. Gusev, C. D. Manning, Customizing an information extraction system to a new domain, in: *ACL2011*, 2011, pp. 2–10.
- [28] H. Zhuge, L. He, Automatic maintenance of category hierarchy, *Future Generation Computer Systems* 67 (2017) 1 – 12.
- [29] ICTY, Prosecutor v. Karadžić: Appeal Judgement, MICT-13-55-0660/2, 2019.
- [30] M. Baumgartner, T. Lampert, Adequate formalization, *Synthese* 164 (2008) 93–115.
- [31] F. Cerutti, G. Pearson, Supporting scientific enquiry with uncertain sources, in: *21st International Conference on Information Fusion*, 2018, pp. 404–411.
- [32] H. Prakken, Argument schemes for discussing Bayesian modellings of complex criminal cases, in: *JURIX 2017*, 2017, pp. 69–78.
- [33] C. S. Vlek, H. Prakken, S. Renooij, B. Verheij, A method for explaining Bayesian networks for legal evidence with scenarios, *Artificial Intelligence and Law* 24 (2016) 285–324.
- [34] B. Verheij, F. Bex, S. T. Timmer, C. S. Vlek, J.-J. C. Meyer, S. Renooij, H. Prakken, Arguments, scenarios and probabilities: connections between three normative frameworks for evidential reasoning, *Law, Probability and Risk* 15 (2016) 35–70.