

Logic of PriCL requires...

truth/falseness (denoted respectively as \top and \perp), conjunction (denoted \wedge), entailment (denoted $A \models B$ if formula A entails formula B), and monotonicity regarding entailment, i.e., if $A \models B$ then $A \wedge C \models B$ for any formula C . As

Definition 1 (Case). A case C is a tuple $(df, CaseDesc, ProofTree, crt)$ s.t.

- df is a formula that we call the decision formula of C .
- $CaseDesc$ is a formula describing the case's circumstances.
- $ProofTree$ is a (finite) tree consisting of formulas f where the formula of the root node is df . Inner nodes are annotated with AND or OR and leaves are annotated with $l \in \{Axiom, Assess\} \cup \{Ref(i) \mid i \in \mathcal{C}_I\}$. Leaf formulas l are additionally associated with a prerequisite formula pre . For leaves annotated with *Axiom*, we require that $pre = l$.
- $crt \in Courts$.

“pre \rightarrow fact” is a description of human judgment, not logical implication

“pres_c” and “facts_c” give all prerequisites/facts in a tree given $CaseDesc \models pre$. Type: $tree \rightarrow set(p/f)$

Subcases are subtrees with around a node n with $df_sub \leftarrow n$

Definition 2 (Subcase). Let $C = (df, CaseDesc, ProofTree, crt)$ be a case and $n \in ProofTree$ a node. Let $sub(n)$ be the subtree of $ProofTree$ with root node n . The case $sub(C, n) := (n, CaseDesc, sub(n), crt)$ is a subcase of C .

df is always necessarily of form $is_legal_action(a)$ because of the nature of privacy law

Definition 3 (Privacy Case). Given world knowledge KB_W and action set $Actions$, a case $C = (df, CaseDesc, ProofTree, crt)$ is a privacy case if $df \in \{\neg is_legal_action(a), is_legal_action(a)\}$ for some action $a \in Actions$, where the is_legal_action predicate is not used in either of KB_W or $CaseDesc$.

Definition 4 (Case Consistency). Let $C = (df, CaseDesc, ProofTree, crt)$ be a case. C is consistent if the following holds (for all nodes n where n_1, \dots, n_k are its child nodes)

- (i) $KB_W \wedge CaseDesc \not\models \perp$
- (ii) $KB_W \wedge CaseDesc \models pres_C$
- (iii) $KB_W \wedge CaseDesc \wedge facts_C \not\models \perp$
- (iv) $\bigwedge_{1 \leq i \leq k} n_i \models n$ if n is an AND step and $\bigvee_{1 \leq i \leq k} n_i \models n$ if n is an OR step

Explain intuition behind OR

Definition 5 (Case Law Database (CLD)). A case law database is a tuple $DB = (\mathbf{C}, \leq_t, \text{must-agree}, \text{may-ref}, \mu, U)$ such that:

- \mathbf{C} is a set of cases. We will also write $C \in DB$ for $C \in \mathbf{C}$.
- $\mu : \mathbf{C} \rightarrow \mathcal{C}_1$ is an injective function such that \mathbf{C} is closed under μ . In the following we will also write $\text{Ref}(D)$ for $\text{Ref}(i)$ if $\mu(D) = i$.
- Let $<_{\text{ref}} := \{(C, D) \mid D \text{ contains a } \text{Ref}(C) \text{ node}\}$ and \leq_t is an order that we call time order of the cases. It has to hold:

$$\begin{array}{c} \text{must-agree} \subseteq \\ <_{\text{ref}} \subseteq \end{array} \text{may-ref} \subseteq \leq_t \subseteq \mathbf{C} \times \mathbf{C}$$

- U specifies the unwarranted nodes, i.e., $U : \mathbf{C} \rightarrow \mathbf{N}$ is function such that
 - \mathbf{N} is a subset of the nodes labelled with **Assess** or **Ref** in the cases \mathbf{C} .
 - The set increases monotonic, i.e., $C \leq_t D \implies U(C) \subseteq U(D)$.

We denote the unwarranted nodes of DB by $U(DB) := \bigcup_{C \in \mathbf{C}} U(C)$.

tuitively, $\text{may-ref}(C_1, C_2)$ denotes the circumstances that case C_1 may reference case C_2 ; $\text{must-agree}(C_1, C_2)$ analogously denotes that C_1 must agree with C_2 .

Explain must-agree vs may-ref / ratio decidendi vs obiter dicta

To be warranted, a case must not require prerequisites

Definition 6 (Warranted Subcase). A subcase $(df, \text{CaseDesc}, \text{ProofTree}, crt)$ is warranted with respect to a set N of nodes if the case $(df, \text{CaseDesc}, \text{ProofTree}', crt)$ is consistent where $\text{ProofTree}'$ is derived from ProofTree by replacing every precondition of a node $n \in N$ by \perp .

A reference is correct if there are shared decision requirements on a node and shared prerequisites

Definition 7 (Correct Case Reference). Let DB be a case law database and $C = (df, \text{CaseDesc}, \text{ProofTree}, crt)$ a case in DB . A leaf node $pre \rightarrow fact$ in ProofTree annotated with $\text{Ref}(D)$ references correctly if $D_u = (fact, \text{CaseDesc}_D, \text{ProofTree}_D, crt_D)$ is a warranted subcase of a case $D \in DB$ w.r.t. $U(C)$, $\text{may-ref}(C, D)$ holds and $KB_W \wedge pre \models pres_D$. C references correctly if all its leaves annotated with $\text{Ref}(D)$ reference correctly.

Cases conflict if their facts or prerequisites contradict and they must agree (because of their df's)

Definition 8 (Case Conflict). Let C_1 be a case in DB and C_2 be a warranted case w.r.t. $U(C_1)$. We say that C_1 is in conflict with C_2 if and only if

- (i) $KB_W \wedge pres_{C_1} \wedge pres_{C_2} \not\models \perp$
- (ii) $KB_W \wedge facts_{C_1} \wedge facts_{C_2} \models \perp$
- (iii) $must\text{-}agree(C_1, C_2)$

A case C is in conflict with DB if there is a $D \in DB$ s.t. C is in conflict with D .

Probably just gloss over this as it closely matches intuition and is way too long:

Definition 9 (Case law database consistency). A case law database $DB = (\mathbf{C}, \leq_t, must\text{-}agree, may\text{-}ref, \mu, U)$ is

- (i) case-wise consistent if every $C \in DB$ is consistent,
- (ii) referentially consistent if every $C \in DB$ references correctly, and
- (iii) hierarchically consistent if every $C \in DB$ is not in conflict with DB .
- (iv) warrants consistently if for every C holds: $U(C)$ contains all $Ref(D)$ nodes where D is an unwarranted subcase w.r.t. $U(C)$.

We call DB consistent if it warrants consistently and is hierarchically, referentially and case-wise consistent.

Deduce = must follow from existing law; Permit = could follow or opposite could follow

Definition 10 (Deducibility and Permissibility). Let $DB = (\mathbf{C}, \leq_t, must\text{-}agree, may\text{-}ref, \mu, U)$ be a consistent CLD, and f a formula. We say that f is permitted in DB under circumstances $CaseDesc$ and court crt if there exists a case $C = (f, CaseDesc, ProofTree, crt)$ such that $ProofTree$ does not contain nodes labeled with *Assess*, and $DB \cup \{C\}$ is consistent (where C is inserted at the end of the timeline \leq_t). We say that f is uncontradicted in DB under $CaseDesc$ and crt if $\neg f$ is not permitted under $CaseDesc$ and crt . We say that f is deducible if it is permitted and uncontradicted.

For sets F of formulas, we say that F is permitted in DB under $CaseDesc$ and crt if there exists a set of cases $\{C_f = (f, CaseDesc, ProofTree_f, crt) \mid f \in F\}$ such that every $ProofTree_f$ does not contain nodes labeled with *Assess*, and $DB \cup \{C_f \mid f \in F\}$ is consistent (where the C_f are inserted in any order at the end of the timeline \leq_t).

This doesn't actually seem that important but does a neat trick in (ii)

Theorem 1. There is a consistent case law database DB , case description $CaseDesc$ and court crt , such that there is a set F of formulas for each of the following properties (in DB under circumstances $CaseDesc$ and court crt):

- (i) For every $f \in F$, f is permissible and F is not permissible.
- (ii) F is permissible, but $\bigwedge_{f \in F} f$ is not permissible.

A supporting set is the facts/pres from which permissibility arises.

Definition 11 (Supporting set). Let $DB = (\mathbf{C}, \leq_t, \text{must-agree}, \text{may-ref}, \mu, U)$ be a consistent case law database, f a formula, CaseDesc a case description and crt a court. A set \mathcal{A} of leaf nodes in DB that are labeled with *Assess* is a supporting set for formula f if the following holds:

- (1) $KB_W \wedge \text{CaseDesc} \models \bigwedge_{(pre \rightarrow fact) \in \mathcal{A}} pre$
- (2) $KB_W \wedge \text{CaseDesc} \wedge \bigwedge_{(pre \rightarrow fact) \in \mathcal{A}} fact \models f$
- (3) $KB_W \wedge \text{CaseDesc} \wedge \bigwedge_{(pre \rightarrow fact) \in \mathcal{A}} fact \not\models \perp$

This trivial follows from definitions but “suggests an algorithm” if you love exponential time

Theorem 2. Let DB be a consistent case law database, f a formula, CaseDesc a case description and crt a court. The following holds:

1. $C \in DB$ with warranted node $f \Rightarrow \exists \mathcal{A}$ that supports f
2. f is permitted (under circumstance CaseDesc and court crt) $\Leftrightarrow \exists \mathcal{A}$ that supports f , is warranted, and is consistent with DB
3. f is deducible $\Leftrightarrow \exists \mathcal{A}$ that supports f and is consistent with DB , and $\forall \mathcal{B}$ it holds that \mathcal{B} does not support $\neg f$, is unwarranted, or is not consistent with DB

This really just clarifies the reasons for the specificity of Thm 2... moving on.

Theorem 3. Let DB be a case law database, and let f be any formula that does not entail \perp . Then there exist cases C_1 and C_2 , each with root node f and the empty case desc \top , such that (inserting C_i at the end of the timeline \leq_t):

- If DB is case-wise consistent, then so is $DB \cup \{C_1\}$.
- If DB is referentially consistent, then so is $DB \cup \{C_2\}$.
- If there is a crt such that $\text{must-agree}(\text{crt}) = \emptyset$, then in addition this holds: for each of $i = 1, 2$, if DB is hierarchically consistent, then so is $DB \cup \{C_i\}$.

I think this last one is pretty obvious.

Theorem 4. The formula \perp is not permitted in any case law database DB , under any circumstances CaseDesc and court crt . The same holds for $\{f, \neg f\}$ if $\text{crt} \in \text{must-agree}(\text{crt})$.

Norms = implicit laws

Definition 12 (Norms). Let $a \in \text{Actions}$. A norm is a formula that has the form $\phi \Rightarrow p$ where $\text{is_legal_action}(a)$ does not occur in ϕ . The norm is a positive norm, denoted ϕ^+ , if $p = \text{is_legal_action}(a)$ and a negative norm, denoted ϕ^- , if $p = \neg \text{is_legal_action}(a)$. A norm ϕ decides p given f if $KB_W \wedge f \models \phi$.

Norms can be found...

Theorem 5. *Let DB be a consistent privacy case law database and $C = (df, CaseDesc, ProofTree, crt) \in DB$. Then there is a norm ϕ that decides df given $CaseDesc$. In particular, there are formulas ϕ_W, ϕ_S such that $is_legal_action(a)$ does not occur in these formulas and*

(1) $facts_C \Rightarrow \phi_W \wedge (\phi_S \Rightarrow df)$ (2) $\phi_W \wedge (\phi_S \Rightarrow df) \Rightarrow df$

And cases restructured around them...

Corollary 1 (Normal forms). *Let $DB = (\mathbf{C}, \leq_t, must-agree, may-ref, \mu, U)$ be a privacy case law database, $C = (df, CaseDesc, ProofTree, crt) \in DB$ be a case, and D be the set of C 's leaf nodes. $N(C)$ is the case that consists of a root node df , two inner nodes ϕ_w and $\phi_S \Rightarrow df$ and the leaf nodes D as children of both inner nodes. We call $N(C)$ the normal form of C . If DB is consistent, then $(\mathbf{C} \setminus \{C\} \cup \{N(C)\}, \leq_t)$ is also consistent (where $N(C)$ is placed at the position of C w.r.t. \leq_t).*

This is determined by a brute force “Algorithm 1: Permissibility” - nevertheless an algorithm though sum^p_2 means “NP with NP oracle”

Theorem 6. *For propositional logic, deciding permissibility is Σ_2^P -complete.*

Theorem 7. *Permissibility is equivalent to satisfiability of a formula whose size is polynomial in the size of DB , $CaseDesc$, and f for*

- (1) *first-order logic.*
- (2) *the description logic \mathcal{ALC} with concept constructors **fills** and **one-of** by role constructors **role-and**, **role-not**, **product**, and **inverse**.²*

The logic to use is “attributive concept language with complements” because of nice properties...