A modal logic formalization of controlled query evaluation

Thomas Studer

Institute of Computer Science
University of Bern
Switzerland

Roadmap

- Controlled query evaluation
- Some logic
- Censor functions
- No-go theorems

Controlled query evaluation (CQE)

Data privacy mechanism where the database (or knowledge base) is equipped with a censor function

Censor checks for each query whether the answer to the query would reveal sensitive information to a user

If this is the case, the censor will distort the answer.

Two options:

- the CQE-system may refuse to answer the query or
- the CQE-system may give an incorrect answer, i.e. it lies

Advantage

This censor based approach has the advantage that the task of maintaining privacy is separated from the task of keeping the data.

This gives more flexibility than an integrated approach (like hiding rows in a database) and guarantees than no information is leaked through otherwise unidentified inference channels.

Applied to a variety of data models:

- complete and incomplete data stores
- relational databases, semi-structured data, ontological knowledge bases

No-go theorems

Well-known in theoretical physics where they describe particular situations that are not physically possible:

- Bell's theorem
- Kochen–Specker theorem
- Frauchiger–Renner paradox

Nurgalieva and del Rio provide a modal logic analysis of the latter paradox.

Arrow's theorem in social choice theory also is a no-go theorem

It states that no voting system can be designed that meets certain given fairness conditions

Pacuit and Yang present a version of independence logic in which Arrow's theorem is derivable.

Our contribution

Develop a highly abstract model for dynamic query evaluation systems like CQE

Formulate several desirable properties of CQE-systems in our framework

Establish two no-go theorems saying that certain combinations of those properties are impossible

Definition: logic

A *logic* L is given by

1. a set of formulas $\mathsf{Fml}_{\mathsf{L}}$ and
2. a consequence relation $\vdash_{\mathsf{L}}$ for L that is a relation between sets of formulas and formulas, i.e. $\vdash_{\mathsf{L}} \subseteq \mathcal{P}(\mathsf{Fml}_{\mathsf{L}}) \times \mathsf{Fml}_{\mathsf{L}}$ satisfying for all $A, C \in \mathsf{Fml}_{\mathsf{L}}$ and $\Gamma, \Delta \in \mathcal{P}(\mathsf{Fml}_{\mathsf{L}})$:
    1. reflexivity: $\{A\} \vdash_{\mathsf{L}} A$;
    2. weakening: $\Gamma \vdash_{\mathsf{L}} A \implies \Gamma, \Delta \vdash_{\mathsf{L}} A$;
    3. transitivity: $\Gamma \vdash_{\mathsf{L}} C$ and $\Delta, C \vdash_{\mathsf{L}} A \implies \Gamma, \Delta \vdash_{\mathsf{L}} A$.

Consistency

1. A logic L is called *consistent* if there exists a formula $A \in \mathsf{Fml}_\mathsf{L}$ such that $\nvdash_\mathsf{L} A$.
2. A set Γ of Fml_L-formulas is called L-*consistent* if there exists a formula $A \in \mathsf{Fml}_\mathsf{L}$ such that $\Gamma \nvdash_\mathsf{L} A$.

Modal logic M over L

The set of formulas $\mathsf{Fml}_{\mathsf{M}}$ is given inductively by:

1. if A is a formula of $\mathsf{Fml}_{\mathsf{L}}$, then $\Box A$ is a formula of $\mathsf{Fml}_{\mathsf{M}}$;
2. $\bot$ is a formula of $\mathsf{Fml}_{\mathsf{M}}$;
3. if A and B are formulas of $\mathsf{Fml}_{\mathsf{M}}$, so is $A \to B$, too.

Semantics

Definition

An M-model $\mathcal{M}$ is a set of sets of Fml_L-formulas, that is

$$\mathcal{M} \subseteq \mathcal{P}(\mathsf{Fml}_\mathsf{L}).$$

Definition

Let $\mathcal{M}$ be an M-model. Truth of an Fml_M-formula in $\mathcal{M}$ is inductively defined by:

1. $\mathcal{M} \Vdash \Box A$ iff $w \vdash_\mathsf{L} A$ for all $w \in \mathcal{M}$;
2. $\mathcal{M} \nVdash \bot$;
3. $\mathcal{M} \Vdash A \to B$ iff $\mathcal{M} \nVdash A$ or $\mathcal{M} \Vdash B$.

Semantics cont'd

Definition

Let Γ be a set of Fml_M-formulas.

1. We write $\mathcal{M} \Vdash \Gamma$ iff $\mathcal{M} \Vdash A$ for each $A \in \Gamma$.
2. Γ is called *satisfiable* iff there exists an M-model $\mathcal{M}$ with $\mathcal{M} \Vdash \Gamma$.
3. Γ *entails* a formula A, in symbols $\Gamma \models A$, iff for each model $\mathcal{M}$ we have that

$$\mathcal{M} \Vdash \Gamma \quad \text{implies} \quad \mathcal{M} \Vdash A.$$

Privacy configuration

A *privacy configuration* is a triple $(\mathsf{KB}, \mathsf{AK}, \mathsf{Sec})$ that consists of:

1. the knowledge base $\mathsf{KB} \subseteq \mathsf{Fml}_\mathsf{L}$, which is only accessible via the censor;
2. the set of a priori knowledge $\mathsf{AK} \subseteq \mathsf{Fml}_\mathsf{M}$, which formalizes general background knowledge known to the attacker and the censor;
3. the set of secrets $\mathsf{Sec} \subseteq \mathsf{Fml}_\mathsf{L}$, which should be protected by the censor.

A privacy configuration $(\mathsf{KB}, \mathsf{AK}, \mathsf{Sec})$ satisfies the following conditions:

1. KB is L-consistent (consistency);
2. $\{\mathsf{KB}\} \Vdash \mathsf{AK}$ (truthful start);
3. $\mathsf{AK} \not\models \Box s$ for each $s \in \mathsf{Sec}$ (hidden secrets).

Queries

A *query* to a knowledge base KB is simply a formula of Fml_L. Given a logic L, we can evaluate a query q over a knowledge base KB. There are two possible answers: t (true) and u (unknown).

Definition

The evaluation function eval is defined by:

$$\mathsf{eval}(\mathsf{KB}, q) := \begin{cases} t & \text{if} \quad \mathsf{KB} \vdash_\mathsf{L} q \\ u & \text{otherwise} \end{cases}$$

Censor

We denote the set of possible answers of a censor by

$$\mathbb{A} := \{t, u, r\}.$$

Let X be a set. Then X^ω denotes the set of infinite sequences of elements of X.

Definition

A *censor* is a mapping that assigns an answering function

$$\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})} : \mathsf{Fml}_{\mathsf{L}}^{\omega} \longrightarrow \mathbb{A}^{\omega}$$

to each privacy configuration $(\mathsf{KB}, \mathsf{AK}, \mathsf{Sec})$. By abuse of notation, we also call the answering function $\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}$ a *censor*. A sequence $q \in \mathsf{Fml}_{\mathsf{L}}^{\omega}$ is called *query sequence*.

Continuous censor

Definition

A censor Cens is *continuous* iff for each privacy configuration (KB, AK, Sec) and for all query sequences $q, q' \in \mathsf{Fml}_{\mathsf{L}}^{\omega}$ and all $n \in \omega$ we have that

$$q|_n = q'|_n \quad \Longrightarrow \quad \mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q)|_n = \mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q')|_n$$

where for an infinite sequence $s = (s_1, s_2, \dots)$, we use $s|_n$ to denote the initial segment of s of length n, i.e. $s|_n = (s_1, \dots, s_n)$.

Continuity means that the answer of a censor to a query does not depend on future queries.

Truthful censor

Definition

A censor Cens is called *truthful* iff for each privacy configuration (KB, AK, Sec), all query sequences $q = (q_1, q_2, \ldots)$, and all sequences

$$(a_1, a_2, \ldots) = \mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q)$$

we have that for all $i \in \omega$

$$a_i = \mathsf{eval}(\mathsf{KB}, q_i) \quad \text{or} \quad a_i = r.$$

A truthful censor may refuse to answer a query in order to protect a secret but it will not give an incorrect answer.

Content

Definition

Given an answer a to a query q, we define its *content* by

$$\mathsf{cont}(q,t) := \Box q \qquad \mathsf{cont}(q,u) := \neg\Box q \qquad \mathsf{cont}(q,r) := \top$$

The content of the answers to a query sequence $q \in \mathsf{Fml}_{\mathsf{L}}^{\omega}$ up to n

$$\mathsf{cont}(\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q), n) := \bigcup_{1 \le i \le n} \{\mathsf{cont}(q_i, a_i)\} \cup \mathsf{AK}$$

where $a = \mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q)$.

Credible censor

Definition

A censor Cens is called *credible* iff for each privacy configuration $(\mathsf{KB}, \mathsf{AK}, \mathsf{Sec})$ and for every query sequence q and every $n \in \omega$, the set $\mathsf{cont}(\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q), n)$ is satisfiable.

A censor is credible if its answers do not contradict each other, that is if they provide a consistent view to an attacker.

Truthful implies credible

Theorem

Every truthful censor is credible.

Effective censor

Definition

A censor Cens is called *effective* iff for each privacy configuration $(\mathsf{KB}, \mathsf{AK}, \mathsf{Sec})$ and for every query sequence $q \in \mathsf{Fml}_{\mathsf{L}}^{\omega}$ and every $n \in \omega$, we have

$$\mathsf{cont}(\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q), n) \not\models \Box s \quad \text{for each } s \in \mathsf{Sec}$$

A censor is effective if it keeps all secrets.

Minimally invasive censor

Definition

Let Cens be an effective and credible censor. This censor is called *minimally invasive* iff for each privacy configuration $(\mathsf{KB}, \mathsf{AK}, \mathsf{Sec})$ and for each query sequence $q \in \mathsf{Fml}_{\mathsf{L}}^{\omega}$, we have that whenever

$$\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q)_i \neq \mathsf{eval}(\mathsf{KB}, q_i),$$

replacing

$$\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q)_i \quad \text{with} \quad \mathsf{eval}(\mathsf{KB}, q_i)$$

would lead to a violation of effectiveness or credibility.

A censor is minimally invasive if it distorts an answer only if otherwise a secret would be leaked.

Repudiating censor

Definition

A censor Cens is called *repudiating* iff for each privacy configuration $(\mathsf{KB}, \mathsf{AK}, \mathsf{Sec})$ and each query sequence q, there are knowledge bases KB_i $(i \in \omega)$ such that

1. $(\mathsf{KB}_i, \mathsf{AK}, \mathsf{Sec})$ is a privacy configuration for each $i \in \omega$;
2. $\mathsf{Cens}_{(\mathsf{KB},\mathsf{AK},\mathsf{Sec})}(q)|_n = \mathsf{Cens}_{(\mathsf{KB}_n,\mathsf{AK},\mathsf{Sec})}|_n$, for each $n \in \omega$;
3. $\mathsf{KB}_i \nvDash_{\mathsf{L}} s$ for each $s \in \mathsf{Sec}$ and each $i \in \omega$.

A repudiating censor can plausibly deny all secrets even if the algorithm of the censor is known to an attacker.

First No-Go Theorem

Theorem

A continuous and truthful censor satisfies at most two of the properties effectiveness, minimal invasion, and repudiation.

Second No-Go Theorem

Definition

A censor is *non-refusing* if it never assigns the answer r to a query.

Theorem

Let L *be based on classical logic. A continuous and non-refusing censor cannot be at the same time effective and minimally invasive.*

Conclusion

Conclusion

Lying is necessary!

Thank you!