Consensus-Finding among Logic-Based Agents

Prof. Éric GRÉGOIRE

June 2016

Scientific Context

Logic-based Artificial Intelligence

Artificial agents with deductive reasoning capabilities

Addressed problem

Computing Forms of Consensuses Among Intelligent Agents

Addressed problem

Computing forms of *consensus* among Intelligent agents

Each agent has

- her own agenda/desires/goals/information
- has full deductive capabilities

A much too limited form of consensus

Intersection of all agendas (and of their logical consequences)

Addressed problem

A more ambitious form of consensus

1. take all agendas together
2. extract a non-contradictory subset of this
3. such that it does not contradict any agent

Such a consensus might contain goals that are not shared by all agents.

However, they can be endorsed by any agent since they do not contradict their own plans.

Addressed problem

Example

Difficult political negotiation to form a government coalition

Each political group has its own objectives and these objectives might be all together conflicting.

Compute one (maximal) subset of all objectives that is not self-contradictory and that does not contradict the plans of any group.

Each group might endorse all these objectives since they do not contradict its own plans…

Addressed problem

Example

3 groups need to find a consensus to form a coalition

Group 1. "Increase Taxation, Do not trim social security. If we do not increase taxation then we do not increase defense spendings"

Group 2. "Trim social security. If we increase taxation then we increase defense spendings"

Group 3. "Do not increase defense spendings"

Can we compute a consensus among these groups?

Logical Preliminaries

(...just what is needed for a basic understanding)

- Boolean variables : a, b, c, d,... *can be true or false*
- Connectives : $\wedge$ (and) $\neg$ (not) $\vee$ (non-exclusive or) $\rightarrow$ (implies)
  remember $a \rightarrow b$ is equivalent to $\neg a \vee b$
- Clauses are disjunction of literals : $a \vee b \vee c \vee \neg d$
- Each formula α can be rewritten as a set (i.e., a conjunction) of clauses (CNF)
- SAT = *is there any truth assignement that satisfies this CNF ?* is NP-complete
- Unsatisfiability is equivalent to logical contradiction $a \wedge \neg a$
- Deduction $\Delta \models \alpha$ is equivalent to $\Delta \cup \{\neg\alpha\}$ is UNSAT
- From any contradiction, we can deduce anything and its contrary !

Consensus (definitions)

$\mathcal{S} = [\Phi_1, \ldots, \Phi_n]$ represents n sources Φ_i where each $\Phi_i \subset \mathcal{L}$ and is satisfiable.

Definition 1. *A set $\Gamma \subset \mathcal{L}$ is a consensus for $\mathcal{S}$ iff $\Gamma \subseteq \bigcup_{i=1}^{n} \Phi_i$ and $\forall\, \Phi_i \in \mathcal{S} : \Gamma \cup \Phi_i$ is satisfiable.*

There always exists at least one consensus, which can be the empty set !

We are interested in maximal consensuses! Two kinds of maximality…

Definition 2. *A consensus Γ for $\mathcal{S}$ is $max_{\subseteq}$ iff $\forall\, \Theta$ s.t. $\Gamma \subset \Theta \subseteq \bigcup_{i=1}^{n} \Phi_i$, $\exists\, \Phi_i \in \mathcal{S}$ s.t. $\Theta \cup \Phi_i$ is unsatisfiable.*
A consensus Γ for $\mathcal{S}$ is $max_{\#}$ iff $\forall\, \Theta$ s.t. $\Gamma' \subseteq \bigcup_{i=1}^{n} \Phi_i$ and $\#\Theta > \#\Gamma$, $\exists\, \Phi_i \in \mathcal{S}$ s.t. $\Theta \cup \Phi_i$ is unsatisfiable.

Consensus (example)

Example (c'ed)

ids = Increase defense spending
it = increase taxation
tss = trim social security

$\Phi_1 = \{ it, \neg it \rightarrow \neg ids\}$
$\Phi_2 = \{ tss, it \rightarrow ids\}$
$\Phi_3 = \{\neg ids\}$

$S = [\Phi_1, \Phi_2, \Phi_3]$
`

There exist 3 $max_{\#}$ consensuses for S :

Consensus1 = $\{\neg it \rightarrow \neg ids, it \rightarrow ids\}$
Consensus2 = $\{\neg ids, \neg it \rightarrow \neg ids\}$
Consensus3 = $\{it, \neg it \rightarrow \neg ids\}$

Can we compute *max consensuses*?

Assume all information is in CNF.

Max consensuses are close to *Maximal Satisfiable Subsets* (*MSSes*) but require additional constraints of satisfiability to be obeyed.

Computing max consensus is as hard as computing MSSes in the worst case...

Computing one $MSS_{\#}$ is in FP^{NP}[wit,log]

Computing one $MSS_{\subseteq}$ is in Opt-T

A practical approach

Good point : SAT and related technologies are often efficient

Re-use and adapt them here!

Basic approach : trim the whole information until a *max* consensus is obtained

Deadlock : we cannot make all satisfiability checks in an iterative manner

How to circumvent this problem?

Main deadlock

1. The agents might have conflicting agendas. We cannot check satisfiability with all of them together !

   We need to test satisfiability with each agent iteratively.

2. If we need to ensure maximality we need to consider every possible ordering among all agents and every sets of clauses to be dropped to ensure satisfiability at each step.

Combinatorial blow-up…

Transformational approach

Re-encode the problem of $max_{\#}$ consensus-finding in such a way that it can be solved using just *one* single discrete optimization procedure…

Transformational approach

Key tool (a variant of) Partial-Max-SAT

Let Σ_1 and Σ_2 be two set of clauses.
Partial-Max-SAT($\Sigma_{1,}$ Σ_2) delivers one maximum cardinality subset of clauses of Σ_1 that are satisfiable together will all clauses of Σ_2.

Σ_1 is called the set of soft constraints.
Σ_2 is called the set of hard constraints.

Example Partial-Max-SAT($U\Phi_l$, Φ_k) delivers one $MSS_\#$ of $U\Phi_i$ that does not contradict Φ_k

But we need to use this tool differently!

Transformational Approach to Multiple Contraction

How to resort to *one single call* to Partial-Max-SAT ?

Whenever one Φ_k is conflicting with $U\Phi_l$, some clauses might need be dropped from $U\Phi_l$.

Roughly

- For each such Φ_k , we create a specific problem using its own variables.
- All problems are linked together so that one single call to Partial-Max-SAT delivers one optimal results in terms of number of clauses to be dropped from $U\Phi_l$.

Transformational Approach

For each Φ_k

$\{(\delta_i \vee \neg\alpha_i)$ s.t . δ_i in $U\Phi_l\}$

Hard constraints

Soft constraints $= \{\alpha_i$ s.t δ_i in $U\Phi_l\}$

Ψ ← Partial-Max-SAT(Σ_1, Σ_2)

$\{\delta_j$ s.t. $(\delta_j$ in $U\Phi_l)$ and $(\alpha_j$ in $\Psi)\}$ is one $\max_{\#}$ consensus for $U\Phi_l$

Experimental Study

Instances

227 instances $\cup\Phi_i$ from planning benchmarks

(translated into Boolean clauses)

all Φ_i are mutally contradictory

Software

MSUnCore (as Partial-Max-SAT solver)

MiniSAT

Camus (for computing MSSes)

Hardware

Intel Xeon E5-2643 (3.30GHz), 8Gb RAM on Linux CentOS.

Time-out 30 min.

Experimental Study

Instance	Γ			Direct Approach			Partial-Max-SAT-based approach					
Name (#Vars #Clauses)	#Var	|Γ|	avg |¬γi|	status	time	avg #CoMSS	#var	#hard	1soft	status	time	#rm
blocks_right_2_p_t5 (406 1903)	67	2	236	memout	?	?	2715	3806	1903	solved	453	43
bomb_b10_t10_p_t1 (1000 1870)	500	2	1757	memout	?	?	3870	3740	1870	solved	2007	4
bomb_b5_t1_p_t2 (240 443)	66	10	78	solved	61	9433	2843	4430	443	solved	0	5
coins_p01_p_t3 (536 1419)	112	10	83	memout	?	?	6779	14190	1419	solved	0	4
coins_p03_p_t2 (368 951)	112	5	157	memout	?	?	2791	4755	951	solved	840	11
coins_p03_p_t5 (872 2355)	112	5	157	memout	?	?	6715	11775	2355	solved	63	10
	112	10	85	memout	?	?	11075	23550	2355	solved	2	6
coins_p05_p_t2 (368 951)	112	5	157	memout	?	?	2791	4755	951	solved	196	9
comm_p02_p_t2 (555 1623)	189	10	140	memout	?	?	7173	16230	1623	solved	5	6
comm_p05_p_t5 (3384 12267)	510	10	366	memout	?	?	46107	122670	12267	solved	72	6
emptyroom_d4_g2_p_t1 (44 130)	32	10	22	solved	922	37875	570	1300	130	solved	0	5
emptyroom_d4_g2_p_t5 (188 586)	32	2	113	memout	?	?	962	1172	586	solved	0	14
emptyroom_d8_g4_p_t3 (244 778)	72	10	51	memout	?	?	3218	7780	778	solved	97	7
ring2_r6_p_t1 (76 215)	54	10	38	memout	?	?	975	2150	215	solved	11	8
ring2_r6_p_t2 (134 402)	54	2	190	memout	?	?	670	804	402	solved	0	26
ring_5_p_t1 (114 242)	70	3	164	memout	?	?	584	726	242	solved	148	12
safe_safe_10_p_t5 (166 357)	21	2	75	solved	635	69924	689	714	357	solved	0	5
safe_safe_30_p_t5 (486 1347)	61	10	43	memout	?	?	6207	13470	1347	solved	44	17
sort_num_s_3_p_t1 (39 106)	27	2	96	memout	?	?	184	212	106	solved	0	10
sort_num_s_3_p_t4 (129 400)	27	10	30	memout	?	?	1690	4000	400	solved	0	8
sort_num_s_4_p_t5 (486 1810)	88	10	62	memout	?	?	6670	18100	1810	solved	3353	10
sort_num_s_6_p_t2 (858 3509)	396	3	925	memout	?	?	6083	10527	3509	memout	?	?
uts_k1_p_t2 (71 204)	25	5	40	solved	337	29328	559	1020	204	solved	0	7
uts_k2_p_t5 (530 1903)	81	10	57	memout	?	?	7203	19030	1903	solved	1114	13
uts_k3_p_t3 (682 2695)	169	10	118	memout	?	?	9515	26950	2695	memout	?	?

Pushing the envelope

The transformational approach has been extended successufelly to handle various expressive extensions

(often using *weighted* Partial Max-SAT)

Pushing the envelope

The transformational approach has been extended (using *weighted* Partial Max-SAT)

to handle preferences

among clauses in each Φ_l

among Φ

stratified information sources

....

and their possible combinations.

Pushing the envelope

Integrity constraints (of various possible forms)

Example : some given clauses in *S* can be required to belong to any consensus

Definition 3. *A set $\Gamma \subset \mathcal{L}$ is a consensus for $\mathcal{S}$ under the constraints Ψ iff $\Gamma \subset \mathcal{L}$ is a consensus for $\mathcal{S}$ and $\forall \alpha \in \Psi : \Gamma \vdash \alpha$.*

Example 2. *In the previous example, $\Gamma = \{it, \neg it \rightarrow \neg ids\}$ is a consensus for $\mathcal{S}$ under the constraint $\Psi = \{it\}$. For example, there is no consensus for $\mathcal{S}$ under the constraint $\Psi = \{\neg tss\}$ since tss is logically conflicting with Φ_2.*

Pushing the envelope

Other preference criterion

Preference for a maximum number of concepts to be completeley agreed on within a consensus:

Let Θ and Ψ be two sets of formulas, we note $\#_{var}(\Theta, \Psi)$ the number of different variables occurring in Θ that are not occurring at all in Ψ.

Definition 3. *A consensus Γ for $\mathcal{S}$ is max$_{\#ac}$ ("ac" standing for agreed concepts) iff for any consensus Γ' for $\mathcal{S}$ s.t. $\Gamma \neq \Gamma'$, we have that* $\#_{var}(\Gamma', \bigcup_{i=1}^{n} \Phi_i \setminus \Gamma') \leq \#_{var}(\Gamma, \bigcup_{i=1}^{n} \Phi_i \setminus \Gamma)$.

Everything that is said in $U\Phi_i$ about the *agreed concepts* is within the consensus.

Pushing the envelope

Experimentations

All instances from the international MUS (Minimal Unsatisfiable Sets) competitions: they are formed of up to 15983000 clauses and 4426000 variables (457459 clauses using 139139 different variables, on average):

MSSes are often a few clauses and so are the $max_{\#}$ consensuses

Each instance was randomly split into $n \in$ [3,5,7,10] mutually conflicting same-size (modulo n) Φ_i

Extremely hard problems!

		$n=3$	$n=5$	$n=7$	$n=10$
1	#solved	235	223	210	207
	time	96	109	119	150
	#var	303643	329599	380110	460194
	#cl	1325632	1855884	2386137	3181517
	$\#\text{cl}_{sol}$	7	2	2	2
2	#solved	117	116	107	102
	time	255	229	238	235
	#var	153553	139909	122367	158878
	#cl	2069215	2599468	3129721	3925100
	$\#\text{cl}_{sol}$	20	40	16	26
3	#solved	290	285	279	266
	time	24	49	77	124
	#var	465177	534802	622374	707358
	#cl	1590761	2121016	2651271	3446653
	$\#\text{cl}_{sol}$	167384	92083	65039	46159
	$\#\text{src}_{sol}$	2	2	2	2
4	#solved	137	135	134	133
	time	57	68	67	71
	#var	30731	37129	43290	52929
	#cl	76711	98629	120547	153423
	$\#\text{cl}_{sol}$	3	2	2	2
5	#solved	232	134	140	135
	time	100	67	71	64
	#var	412784	36274	45688	53290
	#cl	1855884	98629	128933	153423
	$\#\text{cl}_{sol}$	7	2	2	2
6	#solved	121	116	104	100
	time	272	227	239	234
	#var	159659	134720	130672	172960
	#cl	2069215	2599468	3129721	3925100
	$\#\text{cl}_{sol}$	19	39	17	39
7	#solved	211	20	23	20
	time	138	51	83	86
	#var	254986	8706	12264	12752
	#cl	1855884	23560	33337	36649
	$\#\text{cl}_{sol}$	8	2	2	2
8	#solved	60	43	38	35
	time	246	166	176	152
	#var	35809	23867	28892	41552
	#cl	2334344	2864599	3394854	4190236
	$\#\text{cl}_{sol}$	2	2	2	2
	$\#\text{src}_{sol}$	2	2	2	2

Table 1: Experimental Results for 1: $max_{\#}$ 2: $max_{\#ac}$ 3: $max_{\#100\%\Phi_i}$ 4: $max_{\prec}$, 5: $max_{[\Phi_1<\cdots<\Phi_n]}$ 6: $max_{\#}(max_{\#ac})$ 7: $max_{[\Phi_1<\cdots<\Phi_n]}(max_{\#})$ and 8. $max_{\#}(max_{\#ac}(max_{\#100\%\Phi_i}))$.

Pushing the envelope

Extension to a modal logic of necessity and possibility (S5)

$\Box\alpha$ means α is necessary (in all possible worlds)

$\Diamond\alpha$ means α is possible (in some possible world)

We have that $\Box\alpha = \neg \Diamond \neg\alpha$

Example 3. *Let us come back to Example 1 and assume now that agent Φ_3 strengthens her desires and does not want to leave open any possibility in the consensus of having an increase of defense spendings: Φ_3 is now $\{\Box\neg ids\}$ (or equivalently $\{\neg\Diamond ids\}$). There remains only one $max_{\#}$ consensus, namely $\Gamma_2 = \{\Box\neg ids, \neg it \rightarrow \neg ids\}$. Note that $\Box\neg ids$ entails $\neg ids$ in S5. Now, if any Φ_i is then augmented with $\Diamond ids$ then no consensus exists anymore.*

Conclusions

- Logic-based forms of consensues have been proposed.
- Ubiquitous applications in Artificial Intelligence (and other domains).
- Their computation is expected to be hard.
- The transformational method allows maximum consensuses to be computed in an efficient way, very often.

Thank you for your attention!