

Online Situation-Determined Agents and their Supervision

Bitá Banihashemi
York University
Toronto, ON, Canada
bita@cse.yorku.ca

Giuseppe De Giacomo
Sapienza Università di Roma
Roma, Italy
degiacomo@dis.uniroma1.it

Yves Lespérance
York University
Toronto, Canada
lesperan@cse.yorku.ca

Abstract

Agent supervision is a form of control/customization where a supervisor restricts the behavior of an agent to enforce certain requirements, while leaving the agent as much autonomy as possible. In this work, we investigate supervision of an agent that may acquire new knowledge about her environment during execution, for example, by sensing. Thus we consider an agent's *online executions*, where, as she executes the program, at each time point she must make decisions on what to do next based on what her current knowledge is. This is done in a setting based on the situation calculus and a variant of the ConGolog programming language. To reason about such agents, we first define a notion of *online situation-determined agent* which ensures that for any sequence of actions that the agent can perform online, the resulting agent configuration is unique. We then present our formalization of the *online maximally permissive supervisor*.

1 Introduction

In many settings, an agent's behavior needs to be restricted to conform to a set of specifications. For instance, the activities of agents in an organization have to adhere to some business rules and privacy/security protocols. One form of this is *customization*, where a generic process for performing a task or achieving a goal is refined to satisfy a client's constraints or preferences. Process customization includes personalization and configuration and finds applications in number of areas.

A key challenge in such settings is ensuring conformance to specifications while preserving the agent's autonomy. Motivated by this and inspired by supervisory control of discrete event systems (Wonham 2014) De Giacomo, Lespérance and Muise (De Giacomo, Lespérance, and Muise 2012) (DLM) proposed *agent supervision* as a form of control/customization of an agent's behavior. The DLM framework is based on the situation calculus (Reiter 2001) and a variant of the ConGolog (De Giacomo, Lespérance, and Levesque 2000) programming language. DLM represent the agent's possible behaviors as a nondeterministic ConGolog process. It is assumed that some of the agent's actions may be *uncontrollable* (e.g. due to auton-

omy). Another ConGolog process represents the supervision specification, i.e., which behaviors are acceptable/desirable.

DLM formalize a notion of *maximally permissive supervisor* (MPS) that minimally constrains the behavior of the agent in the presence of uncontrollable actions so as to enforce the desired behavioral specifications. The original DLM account of agent supervision assumes that the agent does not acquire new knowledge about her environment while executing. This means that all reasoning is done using the same knowledge base. The resulting executions are said to be *offline executions*.

In this paper, we study how we can apply the DLM framework in the case where the agent may acquire new knowledge while executing, for example through sensing. This means that the knowledge base that the agent uses in her reasoning needs to be updated during the execution. For instance, consider a travel planner agent that needs to book a seat on a certain flight. Only after querying the airline web service offering that flight will the agent know if there are seats available on the flight. Technically, this requires switching from offline executions to *online executions* (De Giacomo and Levesque 1999), which, differently from offline executions, are defined meta-theoretically since at every time point the knowledge base used by the agent to deliberate about the next action is different.

Based on online executions, we define a notion of *online situation-determined agent* which ensures that for any sequence of actions that the agent can perform, the resulting agent configuration is unique. We then present our formalization of the *online maximally permissive supervisor*.

2 Preliminaries

The *situation calculus* (SC) is a well known predicate logic language for representing and reasoning about dynamically changing worlds. Within the language, one can formulate *basic action theories* (BATs) that describe how the world changes as the result of actions; see (Reiter 2001) for details of how these are defined. Hereafter, we will use \mathcal{D} to refer to the BAT under consideration. We assume that there is a *finite number of action types* \mathcal{A} . Moreover, we assume that the terms of object sort are in fact a countably infinite set \mathcal{N} of standard names for which we have the unique name assumption and domain closure. A special predicate $Poss(a, s)$ is used to state that action a is

executable in situation s . The abbreviation $Executable(s)$ means that every action performed in reaching situation s was possible in the situation in which it occurred. We write $do([a_1, a_2, \dots, a_{n-1}, a_n], s)$ as an abbreviation for the situation term $do(a_n, do(a_{n-1}, \dots, do(a_2, do(a_1, s)) \dots))$.

To represent and reason about complex actions or processes obtained by suitably executing atomic actions, various so-called *high-level programming languages* have been defined. Here, we concentrate on (a fragment of) ConGolog that includes the following constructs:

$$\delta ::= \alpha \mid \varphi? \mid \delta_1; \delta_2 \mid \delta_1 \mid \delta_2 \mid \pi x. \delta \mid \delta^* \mid \delta_1 \parallel \delta_2 \mid \delta_1 \& \delta_2$$

In the above, α is an action term, possibly with parameters, and φ is situation-suppressed formula, i.e., a SC formula with all situation arguments in fluents suppressed. As usual, we denote by $\varphi[s]$ the formula obtained from φ by restoring the situation argument s into all fluents in φ . Program $\delta_1 \mid \delta_2$ allows for the nondeterministic choice between programs δ_1 and δ_2 . The intersection/synchronous concurrent execution of programs δ_1 and δ_2 (introduced by DLM) is denoted by $\delta_1 \& \delta_2$. See (De Giacomo, Lespérance, and Levesque 2000) for further details on the remaining constructs.

Formally, the semantics of ConGolog is specified in terms of single-step transitions, using two predicates (De Giacomo, Lespérance, and Levesque 2000): (i) $Trans(\delta, s, \delta', s')$, which holds if one step of program δ in situation s may lead to situation s' with δ' remaining to be executed; and (ii) $Final(\delta, s)$, which holds if program δ may legally terminate in situation s . The definitions of $Trans$ and $Final$ we use are as in (De Giacomo, Lespérance, and Pearce 2010), where the test construct $\varphi?$ does not yield any transition, but is final when satisfied. This results in a *synchronous* test construct which does not allow interleaving (every transition involves the execution of an action). Predicate $Do(\delta, s, s')$ means that program δ , when executed starting in situation s , has as a legal terminating situation s' , and is defined as $Do(\delta, s, s') \doteq \exists \delta'. Trans^*(\delta, s, \delta', s') \wedge Final(\delta', s')$ where $Trans^*$ denotes the reflexive transitive closure of $Trans$.

A ConGolog program δ is *situation-determined* (SD) in a situation s (De Giacomo, Lespérance, and Muise 2012) if for every sequence of transitions, the remaining program is determined by the resulting situation, i.e.,

$$\begin{aligned} SituationDetermined(\delta, s) &\doteq \forall s', \delta', \delta'' \\ Trans^*(\delta, s, \delta', s') \wedge Trans^*(\delta, s, \delta'', s') &\supset \delta' = \delta'', \end{aligned}$$

For example, program $(a; b) \mid (a; c)$ is not SD, while $a; (b \mid c)$ is (assuming the actions involved are always executable). Thus, a (partial) execution of a SD program is uniquely determined by the sequence of actions it has produced. Hence a SD program in a starting situation generates a set/language of action sequences, its executions, and operations like intersection and union become natural.

In the rest, we use \mathcal{C} to denote the axioms defining the ConGolog programming language.

3 Online Situation-Determined Agents

In our account of agent supervision, we want to accommodate agents that can acquire new knowledge about their environment during execution, for example by sensing, and

where their knowledge base is updated with this new knowledge. Thus we consider an agent's *online executions*, where, as she executes the program, at each time point, she makes decisions on what to do next based on what her current knowledge is.

Sensing. A crucial aspect of online executions is that the agent can take advantage of sensing. As in (Lespérance, De Giacomo, and Ozgovde 2008), we model sensing as an ordinary action which queries a sensor, followed by the reporting of a sensor result, in the form of an exogenous action.

Specifically, to sense whether fluent P holds within a program, we use a macro:

$$SenseP \doteq QryIfP; (repValP(1) \mid repValP(0)),$$

where $QryIfP$ is an ordinary action that is always executable and is used to query (i.e., sense) if P holds and $repValP(x)$ is an exogenous action with no effect that informs the agent if P holds through its precondition axiom, which is of the form:

$$Poss(repValP(x), s) \equiv P(s) \wedge x = 1 \vee \neg P(s) \wedge x = 0.$$

Thus, we can understand that $SenseP$ reports value 1 through the execution of $repValP(1)$ if P holds, and 0 through the execution of $repValP(0)$ otherwise.

For example, consider the following agent program:

$$\delta^i = SenseP; [P?; A] \mid [\neg P?; B]$$

and assume the agent does not know if P holds initially. Initially we have $\mathcal{D} \cup \mathcal{C} \models Trans(\delta^i, S_0, \delta', S_1)$ where $S_1 = do(QryIfP, S_0)$ and $\delta' = nil; (repValP(1) \mid repValP(0)); [P?; A] \mid [\neg P?; B]$. At S_1 , the agent knows one of the exogenous actions $repValP(0)$ or $repValP(1)$ can occur, but does not know which. After the occurrence of one of these actions, the agent learns whether P holds. For example, if $repValP(1)$ occurs, the agent's knowledge base is now updated to $\mathcal{D} \cup \mathcal{C} \cup \{Poss(repValP(1), S_1)\}$. With this updated knowledge, she knows which action to do next: $\mathcal{D} \cup \mathcal{C} \cup Poss(repValP(1), S_1) \models Trans(nil; [P?; A] \mid [\neg P?; B], do(repValP(1), S_1), nil, do([repValP(1), A], S_1))$. Notice that with this way of doing sensing, we essentially store the sensing results in the situation (which includes all actions executed so far including the exogenous actions used for sensing). In particular the current KB after having performed the sequence of actions \vec{a} is:

$$\mathcal{D} \cup \mathcal{C} \cup \{Executable(do(\vec{a}, S_0))\}.$$

Note that this approach also handles the agent's acquiring knowledge from an arbitrary exogenous action.

Agent online configurations and transitions. We denote an *agent* by σ , which stands for a pair $\langle \mathcal{D}, \delta^i \rangle$, where δ^i is the initial program of the agent expressed in ConGolog and \mathcal{D} is a BAT that represents the agent's initial knowledge (which may be incomplete). We assume that we have a finite set of primitive action types \mathcal{A} , which is the disjoint union of a set of ordinary primitive action types \mathcal{A}^o and exogenous primitive action types \mathcal{A}^e .

An *agent configuration* is modeled as a pair $\langle \delta, \vec{a} \rangle$, where δ is the remaining program and \vec{a} is the current history, i.e., the sequence of actions performed so far starting from S_0 . The initial configuration c^i is $\langle \delta^i, \epsilon \rangle$, where ϵ is the empty sequence of actions.

The *online transition relation* between agent configurations is (a meta-theoretic) binary relation between configurations defined as follows:

$$\begin{aligned} &\langle \delta, \vec{a} \rangle \rightarrow_{A(\vec{n})} \langle \delta', \vec{a}A(\vec{n}) \rangle \\ &\quad \text{if and only if} \\ &\text{either } A \in \mathcal{A}^o, \vec{n} \in \mathcal{N}^k \text{ and} \\ &\mathcal{D} \cup \mathcal{C} \cup \{Executable(do(\vec{a}, S_0))\} \models \\ &\quad Trans(\delta, do(\vec{a}, S_0), \delta', do(A(\vec{n}), do(\vec{a}, S_0))) \\ &\text{or } A \in \mathcal{A}^e, \vec{n} \in \mathcal{N}^k \text{ and} \\ &\mathcal{D} \cup \mathcal{C} \cup \{Executable(do(\vec{a}, S_0)), \\ &\quad Trans(\delta, do(\vec{a}, S_0), \delta', do(A(\vec{n}), do(\vec{a}, S_0)))\} \text{ is satisfiable.} \end{aligned}$$

Here, $\langle \delta, \vec{a} \rangle \rightarrow_{A(\vec{n})} \langle \delta', \vec{a}A(\vec{n}) \rangle$ means that configuration $\langle \delta, \vec{a} \rangle$ can make a single-step online transition to configuration $\langle \delta', \vec{a}A(\vec{n}) \rangle$ by performing action $A(\vec{n})$. If $A(\vec{n})$ is an ordinary action, the agent must know that the action is executable and know what the remaining program is afterwards. If $A(\vec{n})$ is an exogenous action, the agent need only think that the action may be possible with δ' being the remaining program, i.e., it must be consistent with what she knows that the action is executable and δ' is the remaining program. As part of the transition, the theory is (implicitly) updated in that the new exogenous action $A(\vec{n})$ is added to the action sequence, and $Executable(do([\vec{a}, A(\vec{n})], S_0))$ will be added to the theory when it is queried in later transitions, thus incorporating the fact that $Poss(A(\vec{n}), do(\vec{a}, S_0))$ is now known to hold.

The (meta-theoretic) relation $c \rightarrow_{\vec{a}}^* c'$ is the reflexive-transitive closure of $c \rightarrow_{A(\vec{n})} c'$ and denotes that online configuration c' can be reached from the online configuration c by performing a sequence of online transitions involving the sequence of actions \vec{a} .

We also define a (meta-theoretic) predicate c^\checkmark meaning that the online configuration c is known to be final:

$$\begin{aligned} &\langle \delta, \vec{a} \rangle^\checkmark \text{ if and only if} \\ &\mathcal{D} \cup \mathcal{C} \cup \{Executable(do(\vec{a}, S_0))\} \models Final(\delta, do(\vec{a}, S_0)). \end{aligned}$$

Online situation determined agents. In this paper, we are interested in programs that are SD, i.e., given a program, a situation and an action, we want the remaining program to be determined. However this is not sufficient when considering online executions. We want to ensure that the agent always knows what the remaining program is after any sequence of actions. We say that an agent is *online situation-determined* (online SD) if for any sequence of actions that the agent can perform online, the resulting agent configuration is unique. Formally, an agent $\sigma = \langle \mathcal{D}, \delta^i \rangle$ with initial configuration $c^i = \langle \delta^i, \epsilon \rangle$ is *online SD* if and only if for all sequences of actions \vec{a} , if $c^i \rightarrow_{\vec{a}}^* c'$ and $c^i \rightarrow_{\vec{a}}^* c''$ then $c' = c''$.

We say that an agent $\sigma = \langle \mathcal{D}, \delta^i \rangle$ *always knows the remaining program after an exogenous action* if and only if

$$\begin{aligned} &\text{for all action sequences } \vec{a}, A \in \mathcal{A}^e, \vec{n} \in \mathcal{N}^k \\ &\text{if } \mathcal{D} \cup \mathcal{C} \cup \{Executable(do(\vec{a}, S_0)), \\ &\quad Trans(\delta, do(\vec{a}, S_0), \delta', do([\vec{a}, A(\vec{n})], S_0))\} \text{ is satisfiable,} \end{aligned}$$

$$\begin{aligned} &\text{then there exists a program } \delta' \text{ such that} \\ &\mathcal{D} \cup \mathcal{C} \cup \{Executable(do([\vec{a}, A(\vec{n})], S_0))\} \models \\ &\quad Trans(\delta, do(\vec{a}, S_0), \delta', do([\vec{a}, A(\vec{n})], S_0)). \end{aligned}$$

Essentially, this states that whenever the agent considers it possible that an exogenous action may occur, then she knows what the remaining program is afterwards if it does occur.

We can show that:

Theorem 1 *For any agent $\sigma = \langle \mathcal{D}, \delta^i \rangle$, if δ^i is known to be SD in \mathcal{D} , i.e., $\mathcal{D} \cup \mathcal{C} \models SituationDetermined(\delta^i, S_0)$, and if σ always knows the remaining program after an exogenous action, then σ is online SD.*

Online Runs. For an agent σ that is online SD, online executions can be succinctly represented by *runs* formed by the corresponding sequence of actions. The set $\mathcal{RR}(\sigma)$ of (partial) runs of an online SD agent σ with starting configuration c^i is the sequences of actions that can be produced by executing c^i from S_0 : $\mathcal{RR}(\sigma) = \{\vec{a} \mid \exists c. c^i \rightarrow_{\vec{a}}^* c\}$. A run is *complete* if it reaches a final configuration. Formally we define the set $\mathcal{CR}(\sigma)$ of complete runs as: $\mathcal{CR}(\sigma) = \{\vec{a} \mid \exists c. c^i \rightarrow_{\vec{a}}^* c \wedge c^\checkmark\}$. Finally we say that a run is *good* if it can be extended to a complete run. Formally we define the set $\mathcal{GR}(\sigma)$ of good runs as: $\mathcal{GR}(\sigma) = \{\vec{a} \mid \exists c, c'. \vec{a}.c^i \rightarrow_{\vec{a}}^* c \wedge c \rightarrow_{\vec{a}}^* c' \wedge c'^\checkmark\}$.

4 Online Agent Supervision

DLM's account of agent supervision is based on offline executions and does not accommodate agents that acquire new knowledge during a run. Here, we want to deal with agents that may acquire knowledge through sensing and exogenous actions as they operate and make decisions based on what they know, and we model these as online SD agents.

Assume that we have a (non-deterministic) online SD agent $\sigma = \langle \mathcal{D}, \delta^i \rangle$ whose behavior we want to supervise. Let's also suppose that we have a *supervision specification* δ^s of what behaviors we want to allow in the supervised system and that the system $\langle \mathcal{D}, \delta^s \rangle$ is also online SD. First note that if it is possible to control all the actions of the agent, then the result of supervision can be defined as the intersection of the agent and the specification processes. However in general, some of agent's actions may be *uncontrollable*. These are often the result of interaction of an agent with external resources, or may represent aspects of agent's behavior that must remain autonomous and cannot be controlled directly. Similar to DLM, we model this by the special fluent $A_u(a, s)$ that means action a is uncontrollable in situation s .

We say that a specification δ^s is *online controllable* wrt online SD agent $\sigma = \langle \mathcal{D}, \delta^i \rangle$ iff:

$$\begin{aligned} &\forall \vec{a} a_u. \vec{a} \in \mathcal{GR}(\langle \mathcal{D}, \delta^s \rangle) \text{ and} \\ &\mathcal{D} \cup \{Executable(do(\vec{a}, S_0))\} \not\models \neg A_u(a_u, do(\vec{a}, S_0)) \text{ implies} \\ &\quad (\text{if } \vec{a} a_u \in \mathcal{GR}(\sigma) \text{ then } \vec{a} a_u \in \mathcal{GR}(\langle \mathcal{D}, \delta^s \rangle)). \end{aligned}$$

i.e., if we postfix a good online run \vec{a} for $\langle \mathcal{D}, \delta^s \rangle$ with an action a_u that is not known to be controllable which is good for σ (and so \vec{a} must be good for σ as well), then a_u must also be good for $\langle \mathcal{D}, \delta^s \rangle$. (Note that $\vec{a} a_u \in \mathcal{GR}(\sigma)$ and $\vec{a} a_u \in \mathcal{GR}(\langle \mathcal{D}, \delta^s \rangle)$ together imply that $\vec{a} a_u \in \mathcal{GR}(\langle \mathcal{D}, \delta^i \& \delta^s \rangle)$.) This definition differs from DLM's in that it applies to online runs as opposed to offline runs. Moreover it treats actions

that are not known to be controllable as uncontrollable, thus ensuring that δ^s is controllable in all possible models/worlds compatible with what the agent knows. Note that like DLM, we focus on good runs of the process, assuming that the agent will not perform actions that don't lead to a final configuration of δ^i . The supervisor only ensures that given this, the process always conforms to the specification.

Given this, we can then define the *online maximally permissive supervisor* $m_{ps_{onl}}(\delta^s, \sigma)$ of the online SD agent $\sigma = \langle \mathcal{D}, \delta^i \rangle$ which fulfills the supervision specification δ^s :

$$m_{ps_{onl}}(\delta^s, \sigma) = \mathbf{set}(\bigcup_{E \in \mathcal{E}} E) \text{ where}$$

$$\mathcal{E} = \{E \mid E \subseteq \mathcal{CR}(\langle \mathcal{D}, \delta^i \rangle \& \delta^s)\}$$

and $\mathbf{set}(E)$ is online controllable wrt σ

i.e., the online MPS is the union of all sets of action sequences that are complete online runs of both δ^i and δ^s that are online controllable for the agent σ . Our definition is similar to DLM's, but applies to online runs, and relies on online (as opposed to offline) controllability.

The above definition uses the $\mathbf{set}(E)$ construct introduced by DLM, which is a sort of infinitary nondeterministic branch; it takes an arbitrary set of sequences of actions E and turns it into a program, and has the semantics:

$$\mathit{Trans}(\mathbf{set}(E), s, \delta', s') \equiv \exists a, \vec{a}. a\vec{a} \in E \wedge \mathit{Poss}(a, s) \wedge$$

$$s' = \mathit{do}(a, s) \wedge \delta' = \mathbf{set}(\{\vec{a} \mid a\vec{a} \in E \wedge \mathit{Poss}(a, s)\})$$

$$\mathit{Final}(\mathbf{set}(E), s) \equiv \epsilon \in E$$

For the maximally permissive supervisor $m_{ps_{onl}}(\delta^s, \sigma)$ of the online SD agent $\sigma = \langle \mathcal{D}, \delta^i \rangle$ which fulfills the supervision specification δ^s , where $\langle \mathcal{D}, \delta^s \rangle$ is also online SD, the following properties hold: it always exists and is unique, it is online SD, it is online controllable wrt σ , it is non-blocking (i.e., its execution can always be completed successfully), and it is maximally permissive. See (Banihashemi, De Giacomo, and Lespérance 2016), for proofs and further results on online supervision.

Example 1 Suppose that we have an agent that does not know whether P holds initially, i.e., $\mathcal{D} \not\models P(S_0)$ and $\mathcal{D} \not\models \neg P(S_0)$. Suppose that the agent's initial program is:

$$\delta_1^i = [P?; ((A; (C \mid U)) \mid (B; D))] \mid$$

$$[\neg P?; ((A; D) \mid (B; (C \mid U)))]$$

where all actions are ordinary, always executable, and controllable except for U , which is always uncontrollable. Suppose that the supervision specification is:

$$\delta_1^s = (\pi a. a \neq U?; a)^*$$

i.e., any action except U can be performed. The offline MPS obtained using DLM's definition is different depending on whether P holds: for models of the theory where P holds, the offline MPS is $\mathbf{set}(\{B; D\})$, as the set of complete offline runs of δ_1^s in S_0 is $\{[B; D], [A; C]\}$ and $\mathbf{set}(\{[A; C]\})$ is not controllable wrt δ_1^s in S_0 . For models where P does not hold, the offline MPS is $\mathbf{set}(\{A; D\})$, since the set of complete offline runs of δ_1^s in S_0 is $\{[A; D], [B; C]\}$ and $\mathbf{set}(\{[B; C]\})$ is not controllable wrt δ_1^s in S_0 . As it is not known if P holds, it seems that a correct supervisor should neither allow A nor B . Our definition of online MPS yields the correct result, i.e., $\mathbf{set}(\{\epsilon\})$. \square

As the above example illustrates, DLM have an offline MPS for each model of the theory. Instead, we have a single online MPS that works for all models and can include sensing information when acquired.

Example 2 Supervision can also depend on the information that the agent acquires as it executes. Again, suppose that we have an agent that does not know if P holds initially, and that the agent's initial program is $\delta_2^i = \mathit{Sense}_P; \delta_1^i$.

Again, we have different offline MPSs: in models where P holds, the offline MPS is $\mathbf{set}(\{B; D\})$ and in models where P does not hold the offline MPS is $\mathbf{set}(\{A; D\})$. But since the exogenous report makes the truth value of P known after the first action, we get one online MPS for this agent:

$$m_{ps_{onl}}(\delta_1^s, \langle \mathcal{D}, \delta_2^i \rangle) = \mathbf{set}(\{[QryIfP, repValP(1), B, D],$$

$$[QryIfP, repValP(0), A, D]\})$$

Because the agent queries if P holds, the supervisor has enough information to decide the maximal set of runs from then on in each case. So if the reported value of P is true, then the online supervisor should eliminate the complete run $[A, C]$ as it is not controllable, and if P does not hold, the run $[B, C]$ should be eliminated for the same reason. \square

As well, an action's controllability and whether it satisfies the specification may depend on a condition whose truth only becomes known during the execution. Such cases cannot be handled by DLM's original offline account but our online supervision account does handle them correctly.

In (Banihashemi, De Giacomo, and Lespérance 2016), we define a program construct for execution of an agent as constrained by such an online MPS and a new type of lookahead search construct that ensures nonblockingness.

References

- Banihashemi, B.; De Giacomo, G.; and Lespérance, Y. 2016. Online agent supervision in the situation calculus - Extended version. Technical Report EECS-2016-02, York University.
- De Giacomo, G., and Levesque, H. J. 1999. An incremental interpreter for high-level programs with sensing. In *Logical Foundations for Cognitive Agents: Contributions in Honor of Ray Reiter*. 86–102.
- De Giacomo, G.; Lespérance, Y.; and Levesque, H. J. 2000. ConGolog, a concurrent programming language based on the situation calculus. *Artificial Intelligence* 121(1–2):109–169.
- De Giacomo, G.; Lespérance, Y.; and Muise, C. J. 2012. On supervising agents in situation-determined ConGolog. In *Proc. AAMAS*, 1031–1038.
- De Giacomo, G.; Lespérance, Y.; and Pearce, A. R. 2010. Situation calculus-based programs for representing and reasoning about game structures. In *Proc. KR*, 445–455.
- Lespérance, Y.; De Giacomo, G.; and Ozgovde, A. N. 2008. A model of contingent planning for agent programming languages. In *Proc. AAMAS*, 477–484.
- Reiter, R. 2001. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. The MIT Press.
- Wonham, W. 2014. *Supervisory Control of Discrete-Event Systems*. University of Toronto, 2014 edition.