



MIT Open Access Articles

A Model of Tacit Knowledge and Action

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation	Gal, Y. et al. "A Model of Tacit Knowledge and Action." Computational Science and Engineering, 2009. CSE '09. International Conference on. 2009. 463-468. © 2009 IEEE
As Published	http://dx.doi.org/10.1109/CSE.2009.479
Publisher	Institute of Electrical and Electronics Engineers
Version	Final published version
Citable link	http://hdl.handle.net/1721.1/59323
Terms of Use	Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.

A Model of Tacit Knowledge and Action

Ya'akov Gal
MIT CSAIL, and
Harvard University SEAS

Rajesh Kasturirangan
National Institute of
Advanced Studies
Bangalore, India, and
MIT CSAIL

Avi Pfeffer
Harvard University SEAS

Whitman Richards
MIT CSAIL

I. ABSTRACT

Natural Intelligence is based not only on conscious procedural and declarative knowledge, but also on knowledge that is inferred from observing the actions of others. This knowledge is tacit, in that the process of its acquisition remains unspecified. However, tacit knowledge is an accepted guide of behavior, especially in unfamiliar contexts. In situations where knowledge is lacking, animals act on these beliefs without explicitly reasoning about the world or fully considering the consequences of their actions. This paper provides a computational model of behavior in which tacit knowledge plays a crucial role. We model how knowledge arises from observing different types of agents, each of whom reacts differently to the behaviors of others in an unfamiliar context. Agents' interaction in this context is described using directed graphs. We show how a set of observations guide agents' knowledge and behavior given different states of the world.

II. INTRODUCTION

How do we infer the proper behavior in unfamiliar or sensitive contexts? Suppose you are in a line, waiting to greet a foreign dignitary. Do you bow, or do you shake hands? What is the appropriate protocol? Similarly, at a formal dinner, what would be the proper etiquette — especially in an unfamiliar culture? Even the act of extended eye contact may be a faux pas if you are facing a hostile gang or a malevolent dictator. Studies show that humans (and presumably other animals as well) act appropriately in the world without an explicit representation of all of its features (Brooks, 1991; Gibson, 1977; Bonabeau, Dorigo, & Theraulaz, 1999).

We believe appropriate action must be based on tacit knowledge about how the world works. By tacit knowledge, we mean implicit (though not necessarily unconscious) assumptions about the world that guide action. The mark of tacit knowledge is the lack of deliberation. Tacit knowledge is knowledge whose veracity we do not evaluate using deliberate thought. Our definition of tacit knowledge differs from other uses of this term, for example, by Polanyi (1966), who requires that tacit knowledge be unconscious (perceptual knowledge can be conscious without being evaluable) and by Crimmins (1992), who requires that tacit knowledge involve implicit commitment to beliefs about which one does not have a current opinion.

In this paper, we follow Janik (1990) in defining tacit knowledge as knowledge that in principle can be made conscious and explicit, and evaluable when done so, but is currently

implicit. A simple example clarifies our use of this term. When we choose to run a red light (common in Bangalore, uncommon in Stockholm) we make decisions based on an implicit knowledge of the context and actions appropriate to that context. Such decisions rarely involve deliberation or self-reflexivity (which is the mark of explicit knowledge systems). However, if forced to do so, we may be able to state explicitly the trade-offs involved in running a red light (e.g., the fear of being caught by a policeman vs. the need to get to the destination as quickly as possible) and evaluate the trade-offs. When running a red light, as in many other daily life situations, we act tacitly though we could be fully conscious and deliberative in principle.

Tacit knowledge — in combination with information from other sensory modalities— is often invoked in new contexts. For example, one general class of tacit knowledge revolves around judgments of typicality and mimicry, i.e., “copy someone who looks like he knows what he is doing”. When waiting in line to meet the foreign dignitary, one possible strategy is to follow the actions of someone whose dress and mannerism convey familiarity about the proper etiquette. Here typicality is assessed by looking at dress and mannerisms, both of which are provided by perception. Mimicry and typicality judgments are tacit knowledge since they are not stated formally. We are guided mostly by what other people do, not by rigorous analysis and reasoning (Minsky, 1998).

Tacit knowledge allow agents to acquire knowledge and infer accepted modes of behavior in new contexts without fully reasoning about all the possible consequences of their actions. People use tacit knowledge as a surrogate for what is true in the world. Suppose that a driver that is unfamiliar with the traffic laws in a particular country is waiting to turn right behind a line of cars and can only see a school bus in front. If the school bus turns right on a red light, the driver of the car behind it may choose to follow its actions and do the same. However, this driver may choose not to follow a beat-up sedan. This is because this driver assumes the school bus to be following the local rules, while the driver of the beat-up sedan is believed to be reckless.

This paper presents work towards a computational theory of the way tacit knowledge affects the way agents interact with each other over time. Our approach is inspired by canonical models of perception and language (Marr, 1982; Agre & Rosenschein, 1996; Chomsky, 1965) that have laid down a computational theory (by Marr for perception and Chomsky

for language), i.e., a formal statement of the constraints and conditions regulating a visual or linguistic process.

We consider a setting in which multiple agents need to make decisions, and interact with other agents as defined by a graphical network. We incorporate several underlying principles of tacit knowledge into our model. First, agents use their own actions as a tool for conveying knowledge, and use others' actions as a guide for their own behavior. Second, agents make decisions in a bounded way, that depends, under certain conditions, on the actions of their neighbors. We show how these assumptions facilitate the propagation of knowledge and action in the network, and provide certain guarantees about agents' behavior for different network topologies. Lastly, we use inference mechanisms in the model that are simple in computational complexity, but are sufficient to describe a variety of ways in which agents' accrue knowledge and act in the world. We do not mean to suggest that this model can explain people's performance but rather, that it describes, in a principled, clear way the competence that underlies people's capacity for tacit knowledge.

III. A BASIC MODEL

Initially we will consider a single action with values denoted $+$ (e.g., turning right on red) and $-$. We use notation R_+ and R_- to denote that the action is legal or illegal, respectively. We also introduce a knowledge operator $K_i^t(R_+)$, specifying that action $+$ is known to be legal by agent i at time t or $\overline{K_i^t(R_+)}$, specifying that $+$ is not known to be legal at time t (and similarly for R_-). We drop these super- and sub-script when they are clear from context.

Events are assumed to be consistent, such that R_+ and R_- cannot both hold simultaneously. Knowledge in the model is assumed to be correct, such that $K(R_+) \longrightarrow R_+$. It follows that knowledge is consistent, such that $K(R_+) \longrightarrow \overline{K(R_-)}$. (And similarly for R_-). We also assume that for each agent, one of the following must hold: R_+ is known to be true ($K(R_+)$), R_- is known to be true ($K(R_-)$) or nothing is known ($\overline{K(R_+)}, \overline{K(R_-)}$). Consider the action of turning right on a red light. This means that if an agent knows it is legal to turn right on red, it cannot be the case that the agent knows this action to be illegal. Note that the converse does not hold. For example, $\overline{K(R_-)}$ does not imply that R_- holds. Intuitively, not knowing whether it is illegal to turn right on red does not imply that this action is illegal.

To capture the way different agents make decisions, we introduce the notion of a type. In our model, a type essentially refers to an agent's strategy, and this strategy is specific to each type of agent. Our use of types in this work is distinguished from the traditional usage of this term as representing agents' private information in game theory. Agents' knowledge of the types of agents they interact with is the way our model captures tacit knowledge, as we will soon see.

We now introduce the following four types of agents, representing four different strategies that can occur in this example.

- t_1 (conservative). Choose action $-$. (Never turn right on red)
- t_2 (law abiding). Choose action $+$ if $K(R_+)$. (Turn right on red only if you know it is legal)
- t_3 (risk taker). Choose action $+$ if $\overline{K(R_-)}$. (Turn right as long as you don't know it is illegal)
- t_4 (reckless). Choose action $+$. (Always turn right on red)

A. Interaction Graphs

The relationship between multiple agents' knowledge and their actions is defined by a directed network called an *interaction graph*. Each node in the network represents an agent, an edge (i, j) determines that agent j knows the type of agent i and that agent j can observe the action of agent i . The first clause refers to one of the tenets of tacit knowledge, that of knowing how others behave in various situations, though the manner in which this knowledge is constructed is not explicitly specified. The second clause is commonly used to define interaction within networks. We will denote nodes in the graph by their agent types when it is clear from context.

In the traffic example, the following network represents a possible interaction graph describing a line of cars waiting to turn right, in which an agent of type t_3 is waiting behind an agent of type t_1 who is waiting behind an agent of type t_2 , etc...

$$t_4 \rightarrow t_2 \rightarrow t_1 \rightarrow t_3$$

To formally describe the interaction between tacit knowledge and action, we first detail how knowledge is conveyed for each type of agent. We say that an agent *conveys knowledge* if its actions provide information about rules in the world, or about the knowledge of other agents about rules in the world.

- Types t_1 and t_4 never convey knowledge because the strategies of these types do not depend on their knowledge.
- Type t_3 conveys $K_i(R_-)$ when it is observed to do $-$, and conveys $\overline{K_i(R_-)}$ when it is observed to do $+$.
- Type t_2 conveys $K_i(R_+)$ when it is observed to do $+$ and conveys $\overline{K_i(R_+)}$ when it is observed to do $-$.

Consider for example an agent i of type t_3 (reckless) and an agent j of type t_2 (law abiding). Suppose there is an edge (t_3, t_2) in the interaction graph, and that at time step 1 we have that $K_i^1(R_+), \overline{K_j^1(R_+)}, \overline{K_j^1(R_-)}$ (agent i knows that it is legal to do $+$, and agent j does not know whether it is legal or illegal). Following their type specifications, at time step 1, agent i will do action $+$, and agent j will do action $-$. Now, agent j cannot infer from the actions of agent i that R_+ holds, because a t_3 agent that does $+$ will not convey any knowledge. Therefore at time step 2 we still have that $\overline{K_j^2(R_+)}, \overline{K_j^2(R_-)}$ and following its type specification, agent j will choose to do $-$ again.

We now summarize the rules governing the behavior of any agent in the graph given its directly observed neighbors. Table I lists the actions for a row agent j and column agent i given an edge (i, j) in the interaction graph. Each item in the table is a tuple, in which the left entry states the action

TABLE I
ACTION TAKEN BY ROW AGENT j ON OBSERVING ACTION OF COLUMN AGENT i

	t_1	t_2	t_3	t_4
t_1	$\emptyset (-)$	$- (-)$	$- (-)$	$- (\emptyset)$
t_2	$\emptyset (-)$	$+ (-)$	$- (-)$	$- (\emptyset)$
t_3	$\emptyset (+)$	$+ (+)$	$+ (-)$	$+ (\emptyset)$
t_4	$\emptyset (+)$	$+ (+)$	$+ (+)$	$+ (\emptyset)$

taken by j given that i chooses to do $+$, and the right entry (between parentheticals) states the action taken by j given that i chooses to do $-$. A \emptyset symbol denotes a counter-factual event — an action that cannot be chosen by a given type under the circumstances. For example, it cannot be the case that an agent of type t_1 (conservative) is observed to do action $+$. For example, according to the entry in row t_2 , column t_3 in Table I, when t_3 agent does action $+$, the t_2 agent will do action $-$, as we have shown above.

B. From Knowledge Conditions to Actions

We now show how this knowledge informs agents' decisions in the graph. We state the following theorem that specifies the criteria by which knowledge and action interact in the graph.

Theorem 1. *Let C be an interaction graph. A law-abiding agent j of type t_2 in C will choose to do $+$ at time step $t+l$ if and only if the following hold:*

- There is an agent i in C of type t_2 such that $K_i^t(R_+)$ holds.
- There is a directed path in C from i to j of length $\leq l$ that passes solely through agents of type t_2 .

Similarly, an agent j of type t_3 will choose action $-$ at time $t+l$ if and only if $K_j^t(R_-)$ holds and there is a path from i to j of length l that passes solely through agents of type t_3 .

Using the theorem, we can induce a mapping from any agent i and interaction graph C to a knowledge condition for j , stating that j knows an action is legal ($K_j^{t+l}(R_+)$), j knows the rule is illegal ($K_i^{t+l}(R_-)$), or that i does not know whether the rule is legal or illegal ($\overline{K_j^{t+l}(R_+)}, \overline{K_j^{t+l}(R_-)}$).

This theorem is easy to prove. Take for example a path from R_+ to an agent of type t_2 . Any agent along this path that is not of type t_2 will not convey knowledge of R_+ to t_2 , and therefore t_2 will choose action $-$, as specified by Table I. A corollary to this theorem is that at a given time t , any knowledge that is conveyed by different paths in the interaction graph is consistent. That is, any path of reasoning in the graph will always yield the same action for an agent, according to the theorem.

In our model, although agents can observe their neighbors' types in the graph, they may not be able to infer what their neighbors know. Consider for example two agents i and j , both of type t_3 (reckless). Suppose there is an edge (i, j) in the interaction graph, and that at time t we have that $K_i^t(R_+)$, $\overline{K_j^t(R_+)}, \overline{K_j^t(R_-)}$ (agent i knows that it is legal to do $+$, and agent j does not know whether it is legal to do $+$). At time t , agent i will do $+$. Although agent j will also do action

$+$ at time $t+1$, it will not have learned the rule R_+ because this knowledge cannot be conveyed by the actions of a t_3 type agent.

We now show that the relationships between types and actions can be induced from the knowledge conditions that hold for each agent. The possible knowledge conditions in our examples are as follows:

- 1) $\{K(R_+), \overline{K(R_-)}\}$
- 2) $\{K(R_-), \overline{K(R_+)}\}$
- 3) $\{K(R_+), K(R_-)\}$

Note that the set $\{K_i(R_+), K_i(R_-)\}$ cannot occur because that knowledge is correct. We begin by defining an order over knowledge conditions

$$(2) \succ (3) \succ (1)$$

Intuitively, this order represents a degree of severity: knowing that it is *illegal* to turn right on is considered to be more severe than knowing that it is legal. Similarly, not knowing whether it is illegal to turn right on red is more severe than knowing that this action is legal.

This allows us to reformulate agents' types as a mapping from knowledge states to actions. Thus, an agent of type t_3 (risk taker) chooses action $+$ if (1) and (3) hold, while an agent of type t_2 (law abiding) chooses action $+$ solely if (1) holds. Types t_1 and t_4 choose action $+$ and action $-$ respectively for all possible sets of knowledge predicates. This is shown in Table II. As can be seen in the table, it holds that once an

Type	Knowledge Condition		
	(2)	(3)	(1)
t_1	$-$	$-$	$-$
t_2	$-$	$-$	$+$
t_3	$-$	$+$	$+$
t_4	$+$	$+$	$+$

TABLE II
REFORMULATING TYPES USING KNOWLEDGE CONDITIONS

agent type decides to choose action $+$ for a given knowledge condition, this agent will never choose an action $-$ for a knowledge condition that is more severe. For any of the types shown above, it holds that if a knowledge condition K_1 is more severe than a knowledge condition K_2 , then if an action is known to be legal in K_1 , it is also known to be legal in the other knowledge condition. We call the types whose rules of conduct meet these conditions *monotonic*. As an example of a non-monotonic type, consider a "malicious" agent that chooses action $+$ solely under knowledge condition (2). This agent chooses action $+$ when it knows R_- , and chooses action $-$ when it knows R_+ .

We use the idea of monotonic types to characterize a set of "sensible" types in our domain that facilitate the way knowledge is propagated in the interaction graph. They also serve to limit the space of possible types to consider when performing inference. In general, if there are n binary rules, there are three possible sets of knowledge predicates for each

rule, and the total number of possible sets is thus 3^n . A type is a mapping from each of these sets to an action $+$ or $-$, so the number of possible types is $2^{(3^n)}$, which is doubly exponential in the number of rules. By only considering sensible types, we can reduce this space considerably. By using Theorem 1, one can determine the knowledge and the actions of particular agents without having to enumerate all types.

IV. MULTIPLE ACTIONS

In this section we extend the basic model above to handle multiple actions. Let \mathbf{A} is a set of actions. A value for an action $a \in \mathbf{A}$ specifies whether it is legal (T), illegal (F) or unknown (U). For each agent i , we define a knowledge condition K_i^t to be a function from A to action values. A knowledge condition in this model specifies a value for each action in the domain. For example, if $K_i^t(a) = F$, this means that i knows action a to be illegal at time t . (As before, we drop the subscript when the identity of the agent is clear from context). As before, a type is a mapping from knowledge conditions to actions specifying what actions different agents do given their knowledge about rules in the world.

In the basic traffic example there was a sole action, and thus a knowledge condition described a complete mental state for an agent about the domain. The number of possible mental states is generally exponential in the number of actions, but we can generalize the ideas we introduced in the basic example such that values of particular actions will inform values of other actions. To this end, we first impose a complete ordering over action values that prefers legal actions to unknown actions, and unknown actions to illegal actions. For any action $a \in \mathbf{A}$ we have that

$$a = T \succ a = U \succ a = F$$

We can then define an ordering over actions, such that action a_1 dominates action a_2 if for any agent i at time t , knowledge that a_2 is legal also implies that a_1 is legal. (Note that the direction of the ordering on both sides of the rule is reversed.)

$$a_1 \succeq a_2 \implies K_i^t(a_1) \preceq K_i^t(a_2) \quad (1)$$

We are now ready to define an ordering over knowledge conditions. Let K_i^1, K_i^2 be two knowledge conditions for agent i at two different points in time.

Definition 2. We say that K_i^1 is at least more severe than knowledge condition K_i^2 (denoted $K_i^1 \succeq K_i^2$) if the action value that i knows in K_i^1 dominates the value of the same action in K_i^2 .

$$K_i^1 \succeq K_i^2 \equiv \forall a \in A \quad K_i^1(a) \succeq K_i^2(a)$$

As in the basic traffic example, an order over knowledge conditions implies a notion of severity, but Definition 2 extends this notion to multiple actions. Intuitively, If K_i^1 is more severe than K_i^2 , then if an action a is known to be legal in K_i^2 , then it must be the case that a is known to be legal in K_i^1 . In addition, if a is unknown in K_i^2 , then it cannot be the case that a is known to be legal in K_i^1 .

We illustrate by an extension to the former example in which the speed of the car is a discrete variable. Suppose that the minimal driving speed is 10 MPH, and that speeds can increase at intervals of 10 miles per hour. define an ordering over any two speeds a_1, a_2 such that $a_1 \succ a_2$ if and only if $a_1 > a_2$. We will define an ordering over actions such that driving at a slower speed always dominates driving at a faster speed. Consider two speeds, 50 and 40 MPH. It follows from Equation 1 that if $K_i(50 = T)$ holds for an agent i , then it must be the case that $K_i(40 = T)$. (If 50 MPH is known to be a legal speed, it follows that 40 MPH is also legal.) If $K_i(50 = U)$ holds for an agent, then it must be the case that $K_i(40 = U)$ or that $K_i(40 = T)$. (If it is unknown whether 50 MPH is a legal driving speed, then it cannot be the case that 40 MPH is an illegal driving speed). Lastly, if $K_i(50 = F)$ holds, then any value for $K_i(40)$ is possible. (If 50 MPH is known to an illegal speed, it may be legal or illegal to drive at 40 MPH). In general, for any two actions a_1, a_2 such that $a_1 \succ a_2$, the following knowledge conditions cannot hold according to Definition 1.

$$\begin{aligned} &(K_i(a_1 = U), K_i(a_2 = F)), \\ &(K_i(a_1 = T), K_i(a_2 = U)), \\ &(K_i(a_1 = T), K_i(a_2 = F)) \end{aligned}$$

Referring to our example, consider an agent i that at time step 1 knows that the maximal driving speed is 40 MPH, and at time step 2 learns that the maximal driving speed is 30 MPH. According to Definition 2, we have it that K_i^1 is more severe than K_i^2 . To see this, consider any possible driving speed a . If $a \leq 30$, then it holds for i that $K_i^1(a = T)$ and that $K_i^2(a = T)$. If $a > 40$, then it holds for i that both $K_i^1(a = U)$ and $K_i^2(a = U)$. If $a = 40$, then it holds for i that $K_i^1(a = T)$ and $K_i^2(a = F)$. In all of these cases, we have that $K(a_1) \succeq K(a_2)$, according to Definition 2. Thus, according to our model, the case in which a higher speed is legal is strictly more severe than knowing that a lower speed is legal.

We can use the mechanism above to provide a more general definition of a monotonic type. Recall that a type is mapping from knowledge conditions to actions.

Definition 3. A type T of agent i is monotonic if for any two knowledge conditions K_i^1 and K_i^2 at two different points in time, then the following must hold:

$$K_i^1 \succeq K_i^2 \implies T(K_i^1) \succeq T(K_i^2)$$

The definition of a monotonic type resembles that of a monotonic function. The prescribed action by a type for a more severe knowledge condition must dominate the prescribed action of a type for a less knowledge condition. In effect, this means that a monotonic type will not decrease its speed as it learns more severe information. For example, if an agent type for i prescribes to drive at 30 when $K_i^1(50 = T)$ holds, then it cannot prescribe to drive at 20 when $K_i^2(60 = T)$ because

	LA	RT
LA	$\leq y$	min
RT	$\geq y, \leq z$	$\leq y, \leq z$

TABLE III
LAW ABIDING AND RISK TAKING AGENT TYPES

$K_i^1 \prec K_i^2$ and $30 > 20$. Note that it is quite possible for an agent to break the law while still behaving monotonically. As a simple example, consider an agent j that always drives at a speed that is 10 MPH above the maximal speed that it knows is legal.

Referring again to our example, we define the following types of agents for any two speeds x_1, x_2 and knowledge condition K^1 : A *law abiding* agent will drive at speed x_1 if $K^1(x_1 = T)$ (the agent knows that it is legal to drive at speed x_1). A *risk taker* agent will drive at speed x_1 if $K^1(x_1 > F)$ (the agent will drive at a given speed if it does not know that the legal speed limit is lower). The set of monotonic types we introduced for the binary case generalize to the multi-value case and it can be shown that these types are monotonic, according to Definition 3. A law abiding agent will not increase its speed unless it discovers that the speed limit is higher. A risk taker will never increase its initial speed, it will only lower it when it discovers that the speed limit is lower.

We can extend Theorem 1 to fit the case of multiple actions that consider these two types. A law abiding agent type t will drive at speed y (or below) if there is an agent i in the interaction graph for which $K_i(y = T)$ holds, and a path from this agent to t that is composed solely of law abiding agents. (and similarly for a risk taking agent). Table III lists the actions for a row agent j and column agent i given that an edge (i, j) exists in the graph and that agent i is observed to be driving at speed y . We use notation “RT” to refer to a risk taking agent and “LA” to refer to a law abiding agent. A law abiding agent always conveys that the speed limit is y . Thus a risk taker that is observing a law abiding agent in the network that is driving at speed y will do as follows: if the risk taker is driving at speed below y , it will increase its speed to y , because it learns $K_i(y = T)$, and thus the speed limit is at least y . If the risk taker agent is driving at speed $z > y$, it will not slow down, because it has not learned $K_i(z = U)$. A law abiding agent that observes a risk taker driving at speed y will not change its speed. However, a risk taker that observes another risk taker will increase its speed if it is driving more slowly.

V. DISCUSSION

The interaction graph in our examples above was a simple directed chain, with four different types of nodes arranged in an arbitrary order. There are many other forms of interaction graphs. Obvious cases are a leader of a group (connected graph), or when everyone can see everyone else (complete graph), a bipartite graph, random graphs, etc. With only 8 nodes, there are already 12,000 possible forms (Harary, 1969). Note that Theorem 1, however, is general and lays out a

condition where knowledge of the graph form and node types can lead to a correct action. Theorem 1 illustrates the larger claim that whenever tacit knowledge is governed by underlying principles such as the global consistency of knowledge, agents have enough information to do the right thing. In addition, our use of monotonic types allows to limit the inference to a restricted set of sensible types that facilitate the propagation of knowledge in the network.

The other interesting computational element in our model is the use of monotonicity as a constraint to address a difficult inverse problem, namely, to infer the “correct action” based on observations of other agents behavior. In general, there cannot be a unique solution to the problem of correct action, yet the assumption that people fall into types constrains the inference procedure. However, types alone are not enough. We need to assume that both types of agents and the set of observations are structured in a partial ordering. In our traffic example, types are ordered according to their response to punishment (reckless to conservative). We also assume that agents knowledge is ordered as well, i.e., that no one ever knows less than she knew before so observations always increase knowledge. These two are distinct partial orders but they are related, for whenever there is a scale of punishment, it helps if observations help rather than hurt.

Finally, although the inference process may in some cases be complex, our main intent is to present a simple representation for evaluating how tacit knowledge can lead to correct actions in unfamiliar contexts. The fact that tacit knowledge is largely implicit, suggests how lower animals as well as humans can learn appropriate behaviors without directed examples and tutors.

VI. CONCLUSION AND FUTURE WORK

This work presented a computational model for how tacit knowledge guides actions in settings where agents’ interaction is described using a graphical network. The model incorporates several principles of tacit knowledge that have been mentioned in recent philosophical and psychological theories. First, that agents use others’ actions as a guideline for their own behavior. Second, that they forgo a rigorous analysis and make decisions by only considering local information. We show how these two principles facilitate the propagation of knowledge and action in the network, and provide a set of guarantees about agents’ behavior.

We wish to extend the model in the following ways. First, we wish to capture more general types, such as non-monotonic agents (e.g., an agent that drives at least 45 MPH even if it does not know the legal speed limit, never drives above 65 MPH, and drives at the maximal speed x that it knows to be legal where $45 < x < 65$.) To do so, we will introduce modal logic operators and extend the knowledge conditions to handle probabilistic events. Second, we are incorporating cultural contexts into the model and showing how these effect agents’ reasoning.

REFERENCES

- Agre, P., & Rosenschein, S. (1996). *Computational theories of interaction and agency*. MIT Press.
- Bonabeau, E., Dorigo, M., & Theraulaz, G. (1999). *Swarm intelligence: From natural to artificial systems*. Oxford University Press.
- Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, 47, 139–159.
- Chomsky, N. (1965). *Aspects of the theory of syntax*. MIT Press, Cambridge.
- Crimmins, M. (1992). *Talk about beliefs*. The MIT Press.
- Gibson, J. J. (1977). The theory of affordance. In R. Shaw & J. Bransford (Eds.), *Perceiving, acting, and knowing*.
- Harary, F. (1969). *Graph theory*. Addison-Wesley.
- Janik, A. (1990). *Knowledge, skill and artificial intelligence* (B. Göransson & I. Josefson, Eds.). London: Springer-Verlag.
- Marr, D. (1982). *Vision*. Freeman and Co.
- Minsky, M. (1998). *The society of mind*. Simon and Schuster.
- Polanyi, M. (1966). *The tacit dimension*. London: Routledge & Kegan Paul.