A Modal Logic of Optimality (Student Abstract)

James T. Oswald[1, 2], Brandon Rozek[1, 2], Thomas Macaulay Ferguson[2], Selmer Bringsjord[1, 2]

[1]Rensselaer AI and Reasoning Laboratory
[2]Rensselaer Polytechnic Institute
oswalj@rpi.edu, rozekb@rpi.edu, tferguson@gradcenter.cuny.edu, SelmerBringsjord@gmail.com

Abstract

We present our work on a new modal logic of optimality, $\mathcal{OPT}$, whose semantics are modeled in terms of optimal paths through reward-weighted transition systems. We prove some basic properties of $\mathcal{OPT}$, including its status as a *normal* modal logic, as well as its relation to some of the standard modal axioms. We end with a discussion of applications to AI and future research directions and extensions.

Introduction

The notion of optimality plays a key role in diverse areas of computer science, particularly in AI. Despite this, optimality has yet to be given a formal-logic treatment, even with its clear status as a linguistic modal: the statement "It is optimal that φ" has an intensional truth value. In our semantics we define a formula φ to be optimal with respect to a reward-weighted transition system and a state s iff on all optimal paths p from s (paths which maximize reward through a reward-weighted transition system), φ holds in all states on p past a certain point. Our semantics for optimality are influenced by semantics from branching temporal logics, particularly Computational Tree Logic, CTL (Clarke and Emerson 1981), and its derivatives. In the language of branching temporal logic, our formulation of the optimality of φ is that "φ is eventually always true on all optimal paths." The closest logics in the literature, at least model-wise, are Markov Temporal Logic (MTL_0) (Jamroga 2008) and Discounted Computational Tree Logic (DCTL) (de Alfaro et al. 2005). Both have no specific notions of an optimality operator, and use similar but slightly different models than ours. Our approach places optimality first and foremost as the sole operator, and while influenced by the semantics of CTL, we do not extend it, opting instead for the minimal amount of machinery required to capture a notion of optimality. After defining the syntax and semantics of $\mathcal{OPT}$ in the following section, we prove some basic properties and then turn to a discussion of applications and future work.

Syntax

The syntax of $\mathcal{OPT}$ is that of a standard modal logic, taking $\wedge, \neg$, and $\Box$ as primitive, and assuming all other operators are defined compositionally. Given a set of atomic propositions $\mathbf{At}$, we define formulae φ as follows, where $a \in \mathbf{At}$.

$$\varphi ::= a \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Box\varphi$$

Statements of the form $\Box\varphi$ are read as "It is optimal that φ."

Semantics

$\mathcal{OPT}$ formulae are interpreted over reward-weighted transition systems $\mathcal{M} = (S, T, v, r)$ where:

- S is a finite[1] set of states (possible worlds);
- $T \subseteq S \times S$ is the transition (accessibility) relation;
- $v : \mathbf{At} \to \mathcal{P}(S)$ is a valuation function that assigns atomic propositions to the set of states they hold in; and
- $r : T \to \mathbb{R}$ is a reward function assigning transitions a value.

To define the notion of optimality, we need some additional machinery. First, the notion of a path through a transition system: A path $p : S^{\leq\omega} := (p_0, p_1, p_2, ...)$ is a finite or infinite sequence of states where each successor is reachable by its predecessor via T, formally $\forall 0 \leq i \leq |p| - 2 : T((p_i, p_{i+1}))$.[2] We overload the definition of a reward r to include the total reward of a path p, $r : S^{\leq\omega} \to \mathbb{R} \cup \{+\infty, -\infty\}$, and define it as the sum over the rewards on transitions taken by the path: $r(p) := \sum_{i=0}^{|p|-2} r((p_i, p_{i+1}))$.[3] We define $P : S \to \mathcal{P}(S^{\leq\omega})$ to be the function that maps a state to all paths from that state, $P(s) := \{(p_0, p_1, ...) | s = p_0 \wedge \forall 0 \leq i \leq |p| - 2 : T((p_i, p_{i+1}))\}$. Note that every state s has a "do nothing" path $(s) \in P(s)$, where we define $r((s)) := 0$. We say a path p starting at state s

[1]The requirement that the set of states is finite can be weakened to a frame condition that at least one optimal path exists from each state, which we utilize in our proof that $\mathbf{D}$ holds.

[2]For the sake of intuition, we assume paths may be finite or infinite, but note that the model could easily be modified to represent finite paths as infinite paths, allowing us to exclusively reason over infinite paths.

[3]We leave open the question of dealing with divergence of reward on infinite paths here but provide a stopgap measure. Divergence to $+\infty$ or $-\infty$ is fine, the ability to exploit infinite reward as optimal is coherent, but in the event $r(p)$ diverges such that it is undefined (for example via oscillation) we will say it has a reward of 0 by definition, although this may not be computable.

is optimal iff it has the maximal reward of all paths starting at s. Note that s may have multiple optimal paths. We define the set of optimal paths from a state s formally as $P^*(s) := \arg\max_{p \in P(s)}\{r(p)\}$. Our informal definition of a formula φ being optimal depends on the cardnality of the optimal paths p: (1) If an optimal path is finite, we say φ is optimal on p iff it holds in the terminal state of p. (2) If the optimal path is infinite, we say φ is optimal on p iff there exists a point on p at which φ holds indefinitely. We say a formula φ is optimal in a model $\mathcal{M}$ and state s iff it is optimal on all optimal paths starting at s. With this definition of optimality, we can inductively define the semantic entailment ($\vDash$) of a formula at a state as follows:

$$\mathcal{M}, s \vDash a := s \in v(a)$$
$$\mathcal{M}, s \vDash \neg\varphi := \mathcal{M}, s \nvDash \varphi$$
$$\mathcal{M}, s \vDash \varphi \wedge \psi := (\mathcal{M}, s \vDash \varphi) \wedge (\mathcal{M}, s \vDash \psi)$$
$$\mathcal{M}, s \vDash \Box\varphi := \forall p \in P^*(s) \exists i < |p| \forall j \geq i : (\mathcal{M}, p_j \vDash \varphi)$$

Validity is defined as usual: $\vDash \varphi := \forall \mathcal{M}, s : \mathcal{M}, s \vDash \varphi$.

Properties

Our first task is showing that $\mathcal{OPT}$ is a normal modal logic. A *normal* modal logic is (1) closed under the rule of necessitation and (2) has the **K** axiom hold within it (Chellas 1980, Theorem 4.3).

Lemma 1. $(\vDash \varphi) \Rightarrow (\vDash \Box\varphi)$*, $\mathcal{OPT}$ is closed under the rule of necessitation.*

Proof. If φ is valid then it holds in all states. Optimality requires that φ will hold in a subset of states on all optimal paths. Since we have φ holds in all states, it clearly holds in the required states on all optimal paths. ■

Lemma 2. $\vDash \Box(\varphi \rightarrow \psi) \rightarrow (\Box\varphi \rightarrow \Box\psi)$*, i.e., the* **K** *axiom holds in $\mathcal{OPT}$.*

Proof. After expanding some semantic transformations and generalizing over an arbitrary optimal path p we have that: (1) $\exists i_1 < |p| \forall j_1 \geq i_1 : (\mathcal{M}, p_{j_1} \vDash \varphi) \Rightarrow (\mathcal{M}, p_{j_1} \vDash \psi)$ and (2) $\exists i_2 < |p| \forall j_2 \geq i_2 : (\mathcal{M}, p_{j_2} \vDash \varphi)$ and need to prove $\exists i_3 < |p| \forall j_3 \geq i_3 : (\mathcal{M}, p_{j_3} \vDash \psi)$. Let i_3 be $\max\{i_1, i_2\}$. Then from 1,2 we have that for any $j_3 \geq i_3$, $(\mathcal{M}, p_{j_3} \vDash \varphi) \Rightarrow (\mathcal{M}, p_{j_3} \vDash \psi)$ and $(\mathcal{M}, p_{j_3} \vDash \varphi)$. Apply *modus ponens* and universal introduction to finish. ■

Beyond being a normal modal logic, we want to know what other typical modal axioms hold. We are particularly interested in axioms that correspond with how optimality is used linguistically, particularly: (1) $\mathbf{D}, \Box\varphi \rightarrow \neg\Box\neg\varphi$, read as "If φ is optimal then it is not optimal that not φ," which we expect to be true based on its linguistic usage. (2) $\mathbf{T}, \Box\varphi \rightarrow \varphi$, read as "If φ is optimal then φ," which we expect to be false. (3) $\mathbf{4}, \Box\varphi \rightarrow \Box\Box\varphi$, read as "If φ is optimal then it is optimal that it is optimal that φ," which we expect to be true.

Lemma 3. $\vDash \Box\varphi \rightarrow \neg\Box\neg\varphi$*, i.e.,* **D** *holds in $\mathcal{OPT}$.*

Proof. After expanding semantic transformations we have $\forall p \in P^*(s) \exists i_1 < |p| \forall j_1 \geq i_1 : (\mathcal{M}, p_{j_1} \vDash \varphi)$ and from this must prove $\exists p \in P^*(s) \forall i_2 < |p| \exists j_2 \geq i_2 : (\mathcal{M}, p_{j_2} \vDash \varphi)$. $P^*(s)$ is never empty. Select an arbitrary optimal path p^* as our witness for p. To finish the proof we need to show for all points on p^* there exists a subsequent point j_2 on p^* for which φ holds. Using p^* with our hypothesis gives us a point i_1 which suffices. ■

Lemma 4. $\nvDash \Box\varphi \rightarrow \varphi$*, i.e.,* **T** *does not hold in $\mathcal{OPT}$.*

Proof. Consider a model $\mathcal{M}$ where we have that an atom $a \in$ **At** is false in some state s_1 but true in a terminal state s_2 on optimal paths from s_1. One such $\mathcal{M}$ is $(\{s_1, s_2\}, \{(s_1, s_2)\}, \{(a, \{s_2\})\}, \{((s_1, s_2), 1)\})$ at s_1. Thus we have that, $\mathcal{M}, s_1 \vDash \Box a$ but $\mathcal{M}, s_1 \nvDash a$, *i.e.*, $\mathcal{M}, s_1 \nvDash \Box a \rightarrow a$. ■

Lemma 5. $\vDash \Box\varphi \rightarrow \Box\Box\varphi$*, i.e.,* **4** *holds in $\mathcal{OPT}$.*

Proof. After expanding definitions we have for any path p, $\exists i_1 < |p| \forall j_1 \geq i_1 : (\mathcal{M}, p_{j_1} \vDash \varphi)$ and from this must prove $\exists i_2 < |p| \forall j_2 \geq i_2 \forall q \in P^*(p_{j_2}) \exists i_3 < |q| \forall j_3 \geq i_3 : (\mathcal{M}, q_{j_3} \vDash \varphi)$. Let i_1 witness i_2 and observe that q is just a sub-path of p at/past i_1 on which φ holds. Thus any point j_1 on q can be a witness for i_3, and we can apply the hypothesis to finish the proof. ■

Applications and Future Work

Three potential applications for $\mathcal{OPT}$ are: (1) Application to the theoretical analysis of agents under the Legg-Hutter intelligence measure (Legg and Hutter 2007), in which optimality is tightly coupled with intelligence. (2) Applications to automated planning with action costs, where $\mathcal{OPT}$ could be used to make statements and reason about the optimality of particular states. (3) Applications in machine ethics where obligation could be interpreted as optimality under moral utility.[4] Future-work directions include: (1) Extending the model to handle interpretations over probabilistic transition systems, e.g., investigating labeled Markov chains and Markov decision processes as interpretations. (2) Extending $\mathcal{OPT}$ to be a multi-modal logic over multiple reward functions, allowing us to make complex compound statements about optimality over different objectives. A particularly interesting result of working in this direction would be the ability to represent Pareto optimality in this framework. (3) Bounding the semantics of optimality with respect to a time or resource bound, such as with a discount factor on the rewards, or a hard limit on path lookahead when computing path rewards. (4) Further investigation of the relationship between branching-time logics and $\mathcal{OPT}$, particularly MTL_0. (5) Providing a full axiomatic characterization of $\mathcal{OPT}$, along with a proof theory and the accompanying soundness and completeness proofs. In particular, it is currently hypothesized that the converse of **4**, the $C4$ axiom ($\vDash \Box\Box\varphi \rightarrow \Box\varphi$) holds, as well as the **B** axiom ($\vDash \varphi \rightarrow \Box\neg\Box\neg\varphi$), which when taken with **4**, would give us the **5** axiom ($\neg\Box\neg\varphi \rightarrow \Box\neg\Box\neg\varphi$).

[4]Optimality also figures centrally in the notion — common to economics, the cognitive science of reasoning and decision-making, and technical philosophy of decision-making — of an optimal decision; for a nice introduction to the tri-disciplinary treatment of optimality along this line, see (Johnson and Rips 2015).

References

Chellas, B. F. 1980. *Modal Logic: An Introduction*. Cambridge University Press. ISBN 978-0-51162119-2.

Clarke, E. M.; and Emerson, E. A. 1981. Design and Synthesis of Synchronization Skeletons Using Branching-Time Temporal Logic. volume 131 of *Lecture Notes in Computer Science*, 52–71. Springer.

de Alfaro, L.; Faella, M.; Henzinger, T. A.; Majumdar, R.; and Stoelinga, M. 2005. Model Checking Discounted Temporal Properties. *Theor. Comput. Sci.*, 345(1): 139–170.

Jamroga, W. 2008. A Temporal Logic for Markov Chains. 697–704. IFAAMAS.

Johnson, S.; and Rips, L. 2015. Do the Right Thing: The Assumption of Optimality in Lay Decision Theory and Causal Judgment. *Cognitive Psychology*, 77: 42–76.

Legg, S.; and Hutter, M. 2007. Universal Intelligence: A Definition of Machine Intelligence. *Minds Mach.*, 17(4): 391–444.