Introduction: Artificial Intelligence (AI) Planning algorithms have been successful at complex tasks like enabling a robot to perform simple household chores [1] or play Poker at superhuman levels [2]. However, a common but unstated assumption in each of these cases is that the agent is provided a task-specific model by its designers. For example, the agent intended to perform household tasks is not given a model of Poker or vice versa. Having a human designer specify such a compact model is a reasonable way to build agents of limited scope. However, a more generally-intelligent agent — such as one that is expected to help a human perform the variety of tasks she encounters in daily life — must maintain a richer, *open-scope* model. For any particular task, such a model will contain a large amount of task-irrelevant information. This is detrimental to planning approaches because they must now search through intractably large state and action spaces. Indeed, the average human possesses very large amounts of information about many subjects and objects in the world, but when confronted with a specific task like making coffee, she is able to quickly prune this large open-scope model to focus only on the few objects that are relevant [3]. Thus, I want to address this question through my research: <u>how can we enable agents to autonomously prune task-irrelevant information from open-scope models of their environments?</u>

Intellectual Merit

Proposed Approach: An important consideration for pruning open-scope models is that this process must be more computationally efficient than directly searching through the model's state-action space for a plan. However, this is not the case for many previous approaches, and thus it is more efficient to simply plan directly without explicitly pruning irrelevant states and actions. Recently, I helped introduce a novel approach called *Task Scoping* that, given certain assumptions about the problem, is more efficient than planning. More specifically, task scoping is able to prune irrelevant states and actions from deterministic planning problems expressed in the PDDL language and return a new, reduced PDDL specification that is quicker to plan on and provably preserves all optimal plans to any specified goal. Despite promising results on several planning domains, task scoping is currently applicable only to a limited set of problems because it cannot handle stochastic domains and tends be prohibitively slow when there are a large amount of objects in the model. Thus, I propose to augment task scoping with complementary approaches in order to relax these necessary assumptions on the problem while also preserving task scoping's guarantees and computational advantage over state-space search.

Research Plan: The first extension is to enable task scoping to operate on stochastic domains. A naive way to do this is to apply an 'all outcomes determinization' [4], which would simply treat all stochastic action effects as possible. Thus, any action or object that could possibly affect the goal — no matter how improbable this effect may be — would be considered relevant. For instance, a robotic agent tasked with fetching coffee from the kitchen of a crowded building would reason that every human inside the building is task-relevant because any of them *might* attempt to obstruct the robot. Clearly, taking this approach will likely render task scoping unable to prune much of the state-action space in many open-scope domains. A more fruitful direction is to take inspiration from work in probabilistic planning[5] and prune states and actions that are *probably* irrelevant to the goal and simply redo this pruning if the agent is unable to find a plan. This will break the guarantee that task scoping never prunes states and actions that might be necessary to plan, but should ultimately enable more aggressive pruning with greater computational efficiency. In the aforementioned example, such reasoning could enable the robot to only consider humans in the kitchen as relevant, which would vastly reduce the problem's state-action space.

Another important extension is to enable task scoping to efficiently handle open-scope domains with a large number of objects. The current implementation begins by assuming that only state-variables directly mentioned in the goal are relevant, and then adds other state-variables to the relevant set until a convergence criterion is met. Thus, one way to improve the algorithm's efficiency is to integrate methods like 'SGPlan'[6] that decompose the goal hierarchically into simpler sub-goals for which task scoping could quickly converge upon the relevant set. For instance, the challenging task of cooking a three-course meal for a human could be decomposed into cooking each course independently such that only a small number of ingredients is relevant to preparing each particular course. Another way to reduce the time taken to convergence is to have the algorithm begin with a candidate relevant set that's close to the final relevant set. Recent work like 'PLOI'[7] show that training a Graph Neural Network on simple versions of planning problems enables it to quickly produce such a candidate set of relevant states and objects for much larger planning problems.

To accomplish my research goals and enable relevance to be determined efficiently for a wide range of problems, I will need to bring together insights from a variety of different sub-fields, including AI planning, hierarchical reasoning, machine learning and abstraction. My research will contribute to deeper understanding of the computational complexity and benefits of model-based abstraction, which could help advance other sub-fields of AI. One exciting implication is that this understanding could inform what features models learned by current model-based ML algorithms (like the 'Skills To Symbols' framework [8]) should be present for these models to be useful for the downstream task of planning. This could be a step closer towards a Neurosymbolic AI system that integrates planning and learning. Furthermore, understanding relevance-based abstraction computationally might enable us to better understand how humans discern relevance so effectively.

Broader Impacts While systems that achieve superhuman performance at Poker, Go or video-games are certainly impressive, they have limited scope for practical use. For AI powered robots to be truly impactful to humans around the world, they must be able to operate on an open-scope model containing the large variety of intricate, complex tasks that a human must navigate on a daily basis. Furthermore, in order to safely deploy such robots to share environments with humans - especially vulnerable humans like the elderly or differently-abled - we must be able to provide some guarantees on the behaviours these robots might engage in. My research direction takes an important step towards both of these objectives by enabling AI agents to discern relevance in large open-scope domains while also providing guarantees on the solutions that will be preserved. In the long term, such work could be integral to creating more general AI agents that can assist humans everywhere and serve important societal functions like elderly care.

References [1] C. R. Garrett, T. Lozano-Pérez, and L. P. Kaelbling. Pddlstream: Integrating symbolic planners and blackbox samplers via optimistic adaptive planning, 2020. [2] N. Brown and T. Sandholm. Superhuman ai for multiplayer poker. Science, 365(6456):885–890, 2019. [3] S. Yantis. Goal-directed and stimulus-driven determinants of attentional control, 2000. [4] C. J. Muise, S. A. McIlraith, and C. Beck. Improved non-deterministic planning by exploiting state relevance. Twenty-Second International Conference on Automated Planning and Scheduling, 2012. [5] S. Jiménez, A. Coles, and A. Smith. Planning in probabilistic domains using a deterministic numeric planner. 25th Workshop of the UK Planning and Scheduling Special Interest Group, 2006. [6] Y. Chen, C.-W. Hsu, and B. W. Wah. Sgplan: Subgoal partitioning and resolution in planning. 2004. [7] T. Silver, R. Chitnis, A. Curtis, J. Tenenbaum, T. Lozano-Perez, and L. P. Kaelbling. Planning with learned object importance in large problem instances using graph neural networks, 2020. [8] G. Konidaris, L. P. Kaelbling, and T. Lozano-Perez. From skills to symbols: Learning symbolic representations for abstract high-level planning. Journal of Artificial Intelligence Research, 61:215–289, 2018.