Abstract - *Approved for Public Release*

Most real-world game-theoretic settings, including military ones, are imperfect-information games. A player may be uncertain about the adversary's resources, capabilities, locations, readiness, constraints, *etc*. Algorithms for perfect-information games such as chess or Go do not apply in these settings because they cannot handle uncertainty and associated issues such as deceiving, understanding deception by others, reveal/conceal, and decoys. It is crucial for the nation's defense and economy to achieve and maintain a lead in such strategic reasoning capability. To date, game theory has mainly been used to manually analyze small, stylized models. I propose fundamental research that will enable it to be operationalized for computing optimal strategies for real massive-scale imperfect-information settings. I propose radically new research to make game-theoretic reasoning strong, scalable, and applicable. The research will develop game-independent techniques. It will apply to games with sequential and simultaneous moves, and many of the techniques will apply to continuous-time games also. The work spans from ideas and concepts to mathematical theory to large-scale computational experiments. The prongs include the following.

1) Developing orders of magnitude faster *double oracle algorithms* in several ways.

2) Developing highly scalable algorithms for team games by combining scalable methods for two-player zero-sum games with the best tabular algorithms for team games, which are not scalable.

3) Hybridizing double oracle methods and *imperfect-information subgame-solving techniques*. They may yield the scalability of the former with the superhuman reasoning capability of the latter.

4) Developing drastically better algorithms that iteratively sample the game tree. This includes a) hybridizing the state-of-the-art such algorithm, *ESCHER*, with *extensive-form double oracle algorithms*, and b) designing algorithms that are able to soundly use a sampling distribution that changes during training so computation is not wasted on irrational parts of the tree.

5) Developing the first method that enables planning in a learned abstract search space.

6) Developing a unified mathematical and algorithmic framework that soundly combines double oracle methods, algorithms that iteratively sample the game tree, learned state generalization, subgame solving, and abstract planning in a learned search space.

7) Using the game-theoretic algorithms to significantly advance single-agent reinforcement learning (RL). Framing RL as a game offers a unified perspective on robust RL, imitation learning, inverse RL, and generalization across settings. This has the promise to significantly improve RL algorithms and broaden the scope of game-theoretic techniques.

8) Developing methods for comparing the performance of players - even learning ones - with statistical significance. They will support sound early stopping of experiments to save resources, dynamic extension of experiments if significance has not been reached, and post-hoc inference.

If successful, the proposed research will yield insights, mathematical frameworks, a unified algorithmic framework, and algorithms that will make game-theoretic reasoning significantly stronger, orders of magnitude more scalable, and more applicable to the DoD.

I propose two courses on these topics. I will also include select results into CMU's undergraduate and graduate "Intro to AI" courses. I will give tutorials at leading conferences. I will train top-quality PhD, MS, and BS students via teaching and research mentoring. These include ROTC students. I also interact with the Army AI2C at CMU, further facilitating knowledge transfer. My efforts will contribute highly qualified personnel for the defense and national security workforce.

I plan to visit DoD labs and serve on DoD advisory boards and panels. I also offer unique, powerful ways in which I bring fruits of my research to benefit the DoD.