Introduction: Fit to Target Area

We propose novel research to extend the capabilities of deep reinforcement learning algorithms for encompassing multi-agent, spatio-temporal interactions. Deep learning — learning neural nets with multiple layers — is a scientific breakthrough that has enabled powerful new applications. Deep reinforcement learning (DRL) utilizes deep learning to build agents that act strategically to maximize their performance goals [Minh et al. 2015]. We will research approaches within a DRL framework, that permit novel outcomes including predictive modeling, interpreting opponent strategies, and visualizing, understanding, and explaining the resultant models.

Professional sports constitute a major industry that has high visibility in the media and huge engagement from the public. Fundamentally, the problems of sports analytics are those of multi-agent dynamics in general. With the advancement of high frequency tracking systems and video analysis, more and larger event datasets for sports matches have become available. The growth of sports data has led to the growth of sports data analytics. These developments present a timely opportunity for large-scale machine learning to model complex sports dynamics. Sports analytics has the potential to transform the field of RL by becoming one of its major real-world application areas. Beyond sports, many industries are similarly organized around complex organizations that employ teams to achieve their objectives. Machine learning for extracting value from their data is a top priority for our industry partner SPORTLOGiQ. SPORTLOGiQ is the market leader in providing hockey data from video analysis, which counts most NHL teams among its clients. While SPORTLOGiQ has so far focused on ice hockey, the company has the capacity to exploit techniques for other sports. The knowledge generated will apply to other continuous-flow sports such as soccer and basketball. The team has a strong track record of working with SPORTLOGiQ and each other. Schulte and Bornn are both members of the SFU Sports Analytics Group. Mori and Bornn have published together on trajectory analysis using data from SPORTLOGiQ, as recently as the 2018 MIT Sloan Sports Analytics Conference. Mori and Schulte have published four papers with SPORTLOGiQ researchers as co-authors. This work was supported by NSERC grants (Engage, Strategic Partnership Grants), and has led to three patent applications from SPORTLOGiQ.

Our proposal strongly aligns with the target area Information and Communication Technologies (ICT) and the research topic Advanced Data Management and Analytics.

Fit to Target Area Context: Sports analytics falls within ICT by empowering players, coaches, teams, fans, and the media to access the benefits of sports data at scale. Individual teams can utilize data to improve their competitive performance and their cohesion as a team. Individual players can be empowered to improve their fitness, meet performance targets, and enhance cooperation with teammates. Innovative machine learning will help sports stakeholders to make better decisions and to increase fan engagement.

Fit to Research Topic Advanced Data Management and Analytics:

Management and Analytics of data at scale. The data arrive at high volumes (over 3M events per season), at high velocity (about 3000 events per match), record a variety of heterogeneous events (offensive actions, defensive actions, goals, penalty calls, substitutions,…), and come from a range of data sources, including broadcast video, on-site cameras, and human annotation. The rich structure and heterogeneity of sports data both requires and supports developing new machine learning methods. New methods are required that provide the ability to predict, analyze, and optimize at scale—in real or near-real time. For instance, sports media require near instantaneous analysis of match trends and dynamics. Teams require actionable insights within hours for the next match. The proposed research aims to extract strategically valuable information from the high volume and velocity event information.

Analytics for decision-making. The proposed research will develop analytics tools that allow sports stakeholders to discern opponents' strategies, predict game events, evaluate performance by teams and players, and support their decision-making. This includes descriptive analytics tools for querying and visualizing the data that highlight statistical patterns that are relevant for players, coaches, and the media. Prescriptive analytics tools will provide sports decision makers with personalized recommendations for accurate and effective decisions, to optimize strategic coaching decisions in real-time, and support longer-term plans such as drafting, positioning, and enhancing players.

Section 1: Originality, Quality, Work Plan, Applicants

Objectives

Our *long-term objective* (7-10 years) is to develop RL methods for complex structured multi-agent scenarios. Our *medium-term objective* (3-7 years) is to develop RL methods that solve key problems in sports analytics.

The building blocks of RL models are states, actions, and rewards. Sports are a natural fit for RL models because, essentially, the structure of the model is determined by the rules of the game [Routley and Schulte 2015]. A *state* represents the current match context, including features such as the current score, manpower differentials, game time. A state can also include information about recent match events. SPORTLOGiQ soon expects complete tracking data for the NHL, recording the location of each player at each time. This player location vector would be part of the state. *Actions* are discrete events due to a player, such as shots, blocks, passes, hits etc. SPORTLOGiQ's data record 43 different types of actions-outcome combinations. For each action, their data specify who performs it when and where. *Rewards* in RL are positive events that the agent seeks to maximize. In hockey and soccer, a goal can be taken as the relevant reward for players and teams [Schulte et al. 2017]. The behaviour of a player can statistically be represented as a probability distribution over actions at a game state, which describes what a player tends to do in a given match context. In RL, such a distribution is known as a *policy*. From data such as that provided by SPORTLOGiQ, we can estimate a policy for each player [Schulte et al. 2017, Cervone et al. 2014]. The *value of a state* is the expected reward the agent can expect from its policy given the current game state. Our *short-term objective* (1-3 years) is to develop new approaches for building and leveraging multi-agent RL models to support sports applications.

Objectives for Building RL Models

- *Build hierarchical/multi-level RL models for multi-agent behavioural dynamics*.
- *Develop Agent and Opponent Modeling* to:
  - o Identify types of agents, patterns in action sequences, and tactics. Some important distinctions among agents are not explicitly given as known types, but can be inferred from data. In RL terms, this objective includes methods for *clustering policies*.
  - o Find patterns in agents' strategies, and identify their strengths and weaknesses. As the most interesting patterns are the unusual ones, this objective includes *exception mining*.

Objectives for Leveraging RL Models

- *Develop predictive models for expected game outcomes*. Predictive models enable optimal strategies because they support "what-if" queries to estimate the consequences of a choice on team performance.
- *Investigate strategic optimization and recommendations* for maximizing the success of a group. Our goal is efficient computational methods that analyze and optimize at scale by leveraging predictive RL sports models.
- *Develop an RL-based method for quantifying the performance of an agent/group*. Player and team evaluation is one of the fundamental problems of sports analytics [Schumacher et al. 2010]. Teams invest considerable resources into finding and fielding the right players.

- *Develop interactive visualizations, and analytics tools* for interpreting and explaining RL models and their recommendations to decision makers. In sports, stakeholders include coaches, scouts, players, and analysts.

For each objective, we discuss background, methodology, and milestones. Our interdisciplinary project is at the intersection of two fields: reinforcement learning from computer science and sports analytics from statistics. Our project will develop novel approaches that expand both RL and sports analytics in several new directions. We address the background from each discipline.

Hierarchical/multi-level RL models for multi-agent behavioural dynamics

Overview: In complex structured domains, statistical power is enhanced by describing behavioural dynamics at different levels. In sports, teams and players fall into natural *class hierarchies*: A player may belong to the group of skaters and the subgroup of forwards vs. defensemen. A hierarchical model can learn the behaviour of an average or random forward and how it is different from that of an average defenseman. Another relevant hierarchy natural for sports are compositional *part-whole hierarchies*. For example, teams employ several different lines, and each line is composed of different players. A group hierarchy model can learn the behaviour of different lines. Our goal is to integrate agent hierarchies with behavioural temporal models that model and predict how interactions evolve. Organizing dynamic models according to domain hierarchies makes them easier to interpret for users because the hierarchies reflect their domain knowledge. Hierarchical models also help users understand what individuals have in common and how they are different. For example, a hierarchical model for sports can bring out what different teams, lines, and players have in common and how they are different.

Background: *Reinforcement Learning and Sports Statistics.* RL has been developed extensively over decades [Sutton and Barto 1998]. The key prediction problem in RL is to predict the expected reward for an agent given its current policy. This is different from a standard machine learning setup where the goal is to predict the value of a target/dependent variable from a static vector of features/independent variables: Reinforcement learning models how the effects of an agent's action unfold over time, and takes into account how successive actions interact to determine how much reward the agent achieves. While sports analytics is a successful field [Gudmunsson and Horton 2017], most current predictive models are limited by using the static feature vector format. In contrast, RL models are inherently temporal, and can *directly utilize observation sequences as input*. Some sports analytics papers utilize Markov models to represent temporal dynamics (e.g. [Thomas et al. 2013]). These models use few state features (2-3, e.g. game time, score differential) and no actions. Technically, they represent Markov processes but not Markov *decision* processes, the basis of reinforcement learning.

Behavioural Analytics. Our goal is to *predict, evaluate, and improve* the behavior of real human agents. We thus propose to use RL methods for behavioural analytics for real agents, rather than for the design of artificial agents, which is the goal of most traditional RL applications. Therefore behavioural analytics has been restricted to single-agent domains, or much simpler multi-agent domains than sports (e.g. poker) [Southey et al. 2005]. Two well-established fields of RL that model behavioural dynamics are *behavioural cloning* [Morales and Sammut 2004] and opponent modelling [He et al. 2016]. Behavioural cloning refers to learning an agent's policy from observations of its behavior. In RL terms, our objective is *multi-level behavioural cloning for multi-agent interactions*.

Multi-agent Modelling. Most traditional RL work addresses a single agent. Interest in multi-agent RL has been growing [Buşoniu et al. 2010]. The methods developed in this project will generate new knowledge that extends the theory and applications of multi-agent RL to *complex heterogeneous multi-agent domains with spatio-temporal structure*.

Multi-level Modelling. Integrating agent hierarchies with behavioural models is a virtually unexplored topic in RL. (Hierarchical modeling in RL refers to decomposing a task into subtasks [Kulkarni et al. 2016].) AI researchers have observed that in complex structured domains, *two orthogonal hierarchies* are key to organizing knowledge [Koller and Pfeffer 1997]: an IS-A *class hierarchy* that describes types of objects present in the domain, and a *part-whole hierarchy* that describes how complex objects are composed of constituent objects. Well-known frameworks for class hierarchy models include random effect models [Gelman and Hill 2007], or shrinkage models [Kruschke 2014], which are generally more powerful. A shrinkage model treats probability distributions for higher-level classes as prior distributions for lower-level classes in a Bayesian manner. For example, if we have a model for an entire team, we can use it to define a prior distribution over policies for players from that team, and estimate a posterior distribution for a specific player from this prior and the data. Within sports analytics, shrinkage models have been built for discrete turn-taking sports such as baseball [Kruschke 2014, Ch.9] and cricket [Perera et al. 2018], but not for continuous-flow games such as hockey and soccer. While this work validates the usefulness of hierarchical models for sports, it is based on classical statistical models (linear regression and multinomials), rather than deep temporal neural nets as in our research. Random effects and shrinkage models appear to be a new topic in behavioural cloning. Part-whole hierarchies for group activities represent complex arrangements of players such as lines or formations. Recent work has modelled group behaviour dynamics in space and time [Bialkowski et al. 2014, Yue et al. 2014], with significant contributions from our research team [Ibrahim et al. 2016, Mehrasa et al. 2018]. In RL terms, our work addresses *group behavioural cloning*.

Approach and Methodology: [Lead Researchers Mori and Schulte]. Our main model class are deep neural networks. They support powerful function extrapolation in a nonparametric setting, where a fixed parametric form of the target function is not known. This approach is appropriate for sports data, because it provides sufficient expressive power to capture complex context-dependent dynamics. For behavioural cloning, we will learn a *policy network* that take as input a sequence of game events, and outputs a probability distribution over the player's next action. Sequential data can be modelled with a recurrent neural network (RNN or LSTM). A policy network can also be built for player aggregates at multiple levels. RNN Policy networks have been successfully built for various sports datasets [Mehrasa et al. 2018, Le et al. 2017]. Our research will explore shrinkage models for a *hierarchy of policy networks* for a multi-agent class hierarchy. Current deep temporal models of group activity add parameters for both individual players and a group as a whole. We will build on these models to capture complex group structures, for example lines and formations. The complete model will be a shrinkage model combining class and part-whole hierarchies. SPORTLOGiQ will provide hockey expertise to define useful domain ontologies for hockey and soccer. Temporal models are evaluated in terms of how well they predict future events. SPORTLOGiQ will provide data for real-time updates, and technical personnel to continuously evaluate the predictive accuracy of the new model.

Milestones

- 1 (Year 1-2): Build a class shrinkage policy network model for *player* dynamics. [1 PhD, 1 MSc]
- 3 (Year 1-2): Build a class shrinkage policy network model for *team* dynamics. [1 PhD, 1 MSc]
- 5 (Year 1-2): Build a multi-level policy network model for *group* dynamics, including player lines and formations. [1 Ph.D., 1 MSc]
- 7 (Year 2-3): Combine into a dual hierarchy shrinkage model [1 MSc, with support from Ph.Ds]

Agent/Opponent Modelling: Advanced Behavioural Analytics

Overview: Behavioural Cloning models the observed behaviour of the agents/groups. We will research additional behavioural modelling techniques to include aspects that are not directly observed. These include *clustering* to infer from data additional agent types that are not specified in a known domain hierarchy, and *exception mining* to highlight unusual agents/behaviours.

Background: Clustering has been applied in various areas for sports analytics, for example to identify player types in basketball [Schuhmaker 2010]. Exception mining has been used to highlight exceptional players in the U.K. soccer premier league [Riahi and Schulte 2015]. Most of this work represents players by a single feature vector (e.g. season statistics), not by a dynamic policy. Our project will develop clustering/exception mining for behavioural dynamics, to identify when and in which contexts an agent's behaviour is similar/unusual. Clustering and exception mining for multi-agent dynamics is largely unexplored in both RL and sports analytics.

Approach and Methodology [Lead Researchers Mori and Schulte]: The advanced analytics build on and extend the behavioural cloning shrinkage model.
Clustering. Clustering agents requires a similarity metric. Since our behavioural model represents agents as probabilistic policies, we will apply standard *similarity metrics for distributions*. Previous research suggests that the optimal transport metric is especially well-suited to clustering Markov decision processes [Castro and Precup 2006]. Another option is to extract from the trained policy network a parameter vector that defines an agent, and cluster these vectors.
Exception Mining. A promising framework for our model-based setting is *exceptional model mining* (EMM) [Knobbe et al. 2011, Riahi and Schulte 2015], where an individual is ranked as exceptional to the extent that parameter values characterizing its data differ statistically from those of a random individual. EMM for recurrent neural networks is a new topic.

Milestones

- 2 (Year 3): Develop and evaluate clustering, exception mining *players*. [1 PhD, 1 MSc]
- 4 (Year 3): Develop and evaluate clustering, exception mining for *teams*. [1 PhD, 1 MSc]

Predictive models for expected game outcomes; strategic optimization and recommendation; agent performance evaluation

Overview: Our goal is to develop DRL outcome prediction methods for complex heterogeneous multi-agent domains with spatio-temporal structure. A novel aspect is predicting the expected reward for a group of policies (representing agents), rather than for a single agent interacting with a environment. For example in sports, an important new type of task is predicting the effects of substitutions, which replace one player with another. Predictive models can be applied to recommend optimal strategy choices and to evaluate the performance of agents or groups.

Background: *Value Prediction*. Recurrent networks can be trained to output an estimate of an agent's expected reward, given a sequence of game events. This is called a *value network*. If the input to the network includes a current player action, we have a *Q-value network*. Recent work by a member of our research team [Liu and Schulte 2018] has shown that a Q-value network can be effectively learned from the data provided by SPORTLOGiQ, using well-known techniques such as temporal-difference learning. However, their network made only aggregate predictions for a random home/away team, not for

individual teams and players. While policy networks for group dynamics have been built previously, value networks for groups appear to be a new topic in RL.
Strategic Optimization and Recommendation. Once Q-values have been learned, optimal actions can be extracted by selecting the actions with maximum Q-value (e.g. [Wang et al. 2018]).
Player Evaluation. Recent work by members of our research team [Schulte et al. 2017, Cervone et al. 2014] has shown that the learned Q-values can be used to *evaluate players*: a player's action is scored according to how much it changes the expected reward for his team. Empirical evaluation shows that the Q-value approach performs substantially better than traditional sports statistics that rely on aggregate descriptive statistics (e.g. number of goals scored, plus-minus).

<u>Approach and Methodology</u> [Lead Researchers Mori and Poupart]:
Learning Value Networks. We will investigate scaling Q-value learning for multi-level multi-agent interactions that include both an individual and a group level. A promising approach is to leverage our hierarchical behavioural model to support value learning. One option for this is to train a *joint policy-value network*, which takes as input a match history and outputs both a distribution over next events and expected rewards [Li et al. 2015].
Strategy Optimization. Group-level Q-value networks will allow us to predict expected outcomes for different group states and optimize strategic choices for them (e.g., which formations lead to the highest scoring chances). SPORTLOGiQ will advise on what predictions and strategic recommendations are relevant for sports stakeholders, and provide feedback on the outputs of the algorithms.
Performance Evaluation. SPORTLOGiQ personnel have already included the Q-value approach to player evaluation in their AI platform, based on their patent application. We will apply the Q-value approach to assessing the performance of larger units than players, such as lines and formations. Given that SPORTLOGiQ have dedicated technical personnel, they will take the lead on implementing and evaluating group performance assessment from given Q-values.

<u>Milestones:</u>

- 6 (Year 3): Leveraging the group policy network to learn a group value network [1 Ph.D, 1 MSc.]. Derive and assess strategic recommendations.
- 8 (Year 1-2): Scalable value network learning and real-time strategy optimization. [1 Ph.D., 1 M.Sc.]

Interactive Visualization and Explanation

<u>Overview</u>: To attain maximum impact, machine learning methods must be able to explain the reasons for their recommendations to decision-makers. In sports, the reasons for recommendations and predictions must be explained to coaches, scouts, players, and media analysts.
<u>Background</u>: Much research on interpretability addresses explaining the predictions of neural nets [Boz 2002]; we will build on this work to interpret neural nets that represent policies or value functions for RL, which to our knowledge is an open question. Explaining the predictions and strategic recommendations of continuous-flow multi-agent RL models is another new topic. Sports analysts have often communicated analytics results through visualization. SPORTLOGiQ has developed methods for projecting statistics on to a hockey rink schema, and for linking events of interest to video clips. To visualize dynamical models, Disney research produced an interactive visualization tool for a transition model in basketball, where the user can set up a basketball court scenario and the tool displays which actions the ball holder is likely to perform next [Yue et al. 2014].

<u>Approach and Methodology</u> [Lead Researcher Poupart]:

The Disney research approach can be adapted to let the user explore the Q-values of possible actions in a scenario. We will also investigate using the temporal dynamics model to generate likely future play sequences for a given game state. These play sequences can be combined with the existing NHL simulators to visualize predicted game plays (cf. [Bertasius et al 2018]). We will build on Poupart's previous work for explaining RL policies Khan et al. 2009. SPORTLOGiQ will provide graphics and video expertise to develop the visualizations, and give feedback from their clients on their usefulness. While Poupart will lead on interpretability because of his previous expertise, the team will collaborate on interpreting and visualizing all the research components.

<u>Milestone:</u> 9 (Year 3): Interpreting, explaining and visualizing the predictions and recommendations from the behavioural and value outcome models. [1 Ph.D, 1 MSc.].

Anticipated Impact. Many industries are organized around complex organizations that employ teams to achieve their objectives. This project develops RL as a new framework for modelling, analyzing, and improving the performance of agents and groups in a multi-agent setting. RL models offer a realistic and fine-grained framework for analyzing multi-agent interactions, because they combine a *rich state representation*, which represents a highly informative interaction context, with a model of *temporal dynamics* for the evolution of interactions. The methods developed in this project generate new knowledge that extends the theory and applications of multi-agent RL to *complex heterogeneous multi-agent domains with spatio-temporal structure*. Novel aspects include the following.

1. A complex *real-world environment with massive data*, rather than a low-dimensional synthetic environment with a small dataset.
2. *Groups of agents* that *both cooperate* within their team, and *compete* with other teams, rather than single agents competing with each other.
3. *Predicting the value* or expected reward for a group of policies (representing players) is a new problem. So is predicting the effects of new types of agent-related options, such as substitutions, which replace one player with another.
4. *Advanced Models for Behavioural Analytics,* including multi-level domain hierarchies, clustering, exception mining, and inverse RL.

Our target application domain is sports, with our strong industry partner SPORTLOGiQ. Multi-agent RL offers a fundamentally new approach to the problems of sports analytics that leverages some of the most advanced methods in machine learning and artificial intelligence. At the same time, sports analytics has the potential to transform multi-agent RL by becoming one of its major application areas.

Research Team

We have assembled a team of prominent Canadian researchers with expertise in machine learning for event data, reinforcement learning, and sports analytics.

Dr. Schulte (applicant) has published extensively in leading international venues of machine learning, Artificial Intelligence, and recently, sports analytics. His work has received two best paper awards (IEEE SSCI and Canadian AI). He is a pioneer in applying reinforcement learning to sports analytics. For research on this topic, his student Kurt Routley won the 2017 award for best Master's Thesis from the Canadian Association for Artificial Intelligence. Dr. Schulte will apply his expertise in machine learning and the problems of sports analytics, to develop reinforcement learning for sports analytics.

Dr. Mori (co-applicant) is an international expert on vision-based human activity detection. He serves as Associate Editor for IEEE Transactions on Pattern Analysis and Machine Intelligence (T-PAMI) and on the Editorial Board of International Journal of Computer Vision (IJCV), the top two journals in computer vision. He has published extensively in leading conferences in computer vision and machine learning, including recent papers with the data from SPORTLOGiQ. He won the 2017 Helmholtz Prize (Test of Time Award) for his work on "Recognizing action at a distance." Dr. Mori will apply his expertise on learning temporal models to discover actionable knowledge from dynamic group sports data.

Dr. Poupart (co-applicant) is one of Canada's leading researchers on reinforcement learning. He is well known for developing approximate scalable algorithms for partially observable Markov decision processes. He serves as associate editor of the Journal of Artificial Intelligence Research, and on the editorial board of the Journal of Machine Learning Research. He has won multiple awards, including a best paper award (ICVS), an outstanding collaborator award from Huawei Noah's Ark, an Early Researcher Award from the Ontario Ministry of Research and Innovation (2008), two Google reserch awards. His recent research has addressed hockey analytics. Dr. Poupart will contribute his expertise on applying reinforcement learning concepts in real-world problems, and on scaling reinforcement learning algorithms to the volume, variety, and velocity of sports data. Dr. Poupart is a Professor at the University of Waterloo, well-placed to collaborate with SPORTLOGiQ's new Waterloo research site.

Luke Bornn (collaborator) is an expert on statistical methods for sports analytics and has published extensively in the leading venues, such as the MIT Sloan conference. Dr. Bornn has a unique perspective as a high-profile academic (faculty in the SFU statistics department, and previous member of Harvard's world-leading sports analytics group) who also has a senior executive position in a sports team (VP Strategy and Analytics with the NBA Sacramento Kings). He will contribute his expertise on developing statistics for sports applications, including basketball, soccer, and continuous-flow sports in general.

Work Plan

Coordination. The PI Dr. Schulte will be responsible for overall project work plan coordination. Since the project has multiple locations, a large number of students, and multifaceted collaboration and technology transfer and outreach plans, our budget includes funding for a project manager to assist him. Software repositories for the project are already in place from previous collaborations. Version-controlled software and access-controlled data will be available to the entire project team.

Collaboration. The research team is based in Vancouver, Waterloo, and Montreal. SPORTLOGiQ will provide data, market direction, and results validation. The company has been closely involved in the development of this proposal, and will remain involved in all stages of the project. The research team will conduct weekly teleconferences to coordinate between the Waterloo and Vancouver sites. SFU and Waterloo students will go on exchange visits to allow transfer of ideas between institutions and joint work on project activities.

Equipment. The team's labs at Waterloo, SFU are well equipped with enough workstations to provide one for each student, a total of 15 GPUs, and 3 database servers. In addition the researchers have access to and experience with the extensive resources of Compute Canada.

Project Flow. Our methodology shows lead researchers and graduate students assigned to the separate objectives and milestones. Each co-applicant will be responsible for managing the students in their group. While the primary responsibility for each milestone is assigned to a single co-applicant, the researchers will share their complementary expertise for the different objectives. For example, Schulte's group has built Q-value networks for player evaluation, and will collaborate with Poupart's on the strategic recommendation milestone. Poupart's group has developed Bayesian shrinkage models for sequential data ("Online Bayesian Transfer Learning for Sequential Data Modeling") and will collaborate on this milestone with Schulte's group. Bornn will share his expertise from previous work, which is relevant to all milestones.

Project Scale. Our budget supports a substantial team of researchers and students. This is required for success in the sports analytics industry for several reasons. (1) To have impact with industry stakeholders, a project needs to cover a range of aspects, such as predictive accuracy, domain knowledge, explaining results, and customizability. Tackling only one aspect may be sufficient for a research paper but not to impact real-world decision-making. (2) Sports analytics is a fast-moving field, due to contributions from major U.S. and European research groups with strong well-resourced collaborators in universities and industry (e.g. STATS LLC). Canada has a lead in hockey analytics, but to maintain and expand it requires the resources of a Strategic Project Partnership.

Section 2: Training Potential and Training Plan

All of the research will provide unique opportunities for student training with direct involvement from the entire project team (Mori, Poupart, Schulte, Bornn, SPORTLOGiQ).

Supervision. The research work will be carried out as a connected series of 4 PhD and 9 MSc theses. Each graduate student will receive extensive expert mentoring, through their direct supervisor and through the collaborations of the different research groups. While ad-hoc teleconferences and remote meetings will form the bulk of these interactions, we have included substantial budget for collaboration travel. This will permit travel for the students and PIs to collaborate closely and provide a strong training environment with sports experts from SPORTLOGiQ.

Industrial Training. SFU students will spend semesters at SPORTLOGiQ's head office in Montreal as interns; we plan for three internships a year. Waterloo students will work regularly with the researchers at the SPORTLOGiQ Waterloo research site. The internships will provide the company privileged access to talented data scientists who have been trained to work with sports data. SPORTLOGiQ personnel will be involved in training directly when through the internships, and indirectly through guidance and feedback at remote and in-person meetings. Their involvement will help the students learn about the real-world technical and business requirements of sports analytics.

Skills Acquired. The proposed research requires students to learn a variety of skills relevant to research and employment in computer science. Through the collaboration students and company personnel will acquire high-demand skills in statistical methods, mathematical modeling, and optimization, that can be applied in many application areas outside sports analytics – e.g. data mining, marketing, and financial industries. These data science skills are in demand with many Canadian organizations.

HQP Track Record. As the team CVs show, the team has a strong track record in research training. Each co-applicant has supervised over 50 HQP. Our students have interned at international (Google, Microsoft Research, Disney Research) and Canadian companies (Big Park, D-Wave Systems), performing work that is directly relevant to their training. Our former students have all gone on to further study or employment in areas relevant to their training. On the academic side, these include professors (Boyd: Calgary; Hoey: Waterloo; Leonard: Johns Hopkins; Wang: Manitoba; Khosravi: Queensland), and many BSc/MSc students pursuing graduate degrees at excellent research schools (Harvard, Princeton, Stanford, Toronto, UBC, Caltech, Georgia Tech,..). On the industry side, alumni have taken many leading roles related to their studies (e.g. Yang, Google Research; Elinas, senior research scientist at Canon; Murray, founder Point Grey Research).

Section 3: Interactions With Supporting Organization

Past Interactions. The team has a strong track record of working with SPORTLOGiQ, including four joint publications and three patent applications, supported by NSERC programs (Engage, SPG-P). The company has been closely involved in the development of this proposal since January 2017, and covered Dr. Schulte's travel expenses for in-person consultation in Montreal to work on the proposal.

Team meetings: Company personnel will be available for bi-weekly teleconferences. In addition there will be in-person strategy meetings (held at project kick-off, 1y, 2y, 3y, plus additional visits to SPORTLOGiQ): These in-person meetings will involve all applicants and SPORTLOGiQ. There will be 3 meetings in Montreal and one in Waterloo. Students will participate by teleconference in the Montreal meetings; Waterloo students will attend the Waterloo meeting. At each meeting, the researchers will present their research directions and recent results. Sportlogiq will describe current technology status and market needs. At the year-end meetings, opportunities for knowledge/technology transfer and intellectual property licensing will be identified. Mehrsan Javan, SPORTLOGiQ's CTO, has visited Vancouver at least once a year over the last three years as part of our on-going collaboration. He will continue the annual trips to Vancouver, and add consultation visits with the Waterloo team at SPORTLOGiQ's research site.

Outreach: The year-end meetings will include an open sports analytics workshop for the local research communities in Waterloo and Montreal. The project manager will assist with the organization of these events. Both Waterloo and Montreal offer great resources for reinforcement learning; for example Microsoft Research, Deep Mind, and the Montreal Institute for Learning Algorithms host many researchers with a focus on RL. The sports analytics workshops will introduce SPORTLOGiQ to the research communities and support their recruiting, and provide the research team and their students with an opportunity to present their work to other experts in RL and machine learning.

Technology Transfer: As shown in the CVs, each researcher has extensive experience with technology transfer, often supported through NSERC and MITACS programs, working with startups, mid-sized tech companies such as RBC Borealis, and industry giants such as Huawei, Google, and Intel. SPORTLOGiQ is the market leader in hockey data. They were one of the official sponsors of this year's MIT Sloan Sports Analytics Conference, the most visible venue in sports analytics. Their clients include the majority of the NHL teams. One of the company's main objectives is to expand their value chain by complementing their video-based data collection with state-of-the-art analytics products for their clients. They have hired a substantial number of full-time personnel to work on analytics, about 10 FTE. A current major direction of expansion is to add data collection and analytics support for other sports, especially soccer. SPORTLOGiQ already has an AI platform in place that incorporates the IP generated by the previous collaborations. New research results can therefore be integrated continuously, becoming ready for customer use within an anticipated timeframe of 1-3 years from project-end. Both SPORTLOGiQ and the researchers have a strong track record of communicating sports analytics results, as shown in their CVs. The company's work has been featured in news segments and broadcast TV. They have publicized applications, case studies, and demos, resulting from this project, in media news releases and social media (Twitter, Youtube), and will do so for the project results. Researchers will continue to communicate results at conferences, workshops, and seminars, both general AI/machine learning/computer vision venues, and specialized sports analytics conferences (Sloan, Vancouver hockey analytics).

Intellectual Property: SportLogiq will have right of first refusal to exclusively license for commercial exploitation of all IP generated. Any IP agreement will be in accordance with NSERC and SFU's policies.

Section 4: Benefits to Canada and Supporting Organization

The sports industry is a huge market in North America with an estimated size of $485 billion US. Given the industry size and the increased attention to the use of statistics in sports, *sport analytics* has become a fast growing industry. Sport analytics has attracted much attention in the research community. As an example, a recent study indicates that the number of Sport analytics related articles published in the past seven years is four times the number of applied operations research articles published in the same time frame [Cochran 2010]. According to Businesswire, by 2022 the sports analytics market is expected to grow by 40% to $3.97B, compared to 28% growth for social media analytics to $5.4B. A strategic investment in sports analytics will allow Canadian researchers to remain innovation leaders in the fields of machine learning, AI, and data science, and create economic opportunities for Canadian companies in sports analytics. This project will seed future industry-university collaborations in multi-agent RL. As a market leader in hockey data, SPORTLOGiQ can leverage connections to NHL teams and the league to exploit the results within an anticipated timeframe of 1-3 years, advancing the company revenue by an estimated 25%. The proposed research will generate actionable analytics that benefit the entire Canadian hockey industry (teams, media, players, player agencies). Further, the results of this collaboration are not limited to sports and the developed algorithms are transferable to other multi-agent domains that involve team performance.

Addressable Markets: Across the five top leagues (Premier League Football (soccer), the NFL (football), NBA (basketball), MLB (baseball), and NHL (hockey)), teams are currently spending approximately $400M annually on advanced analytics hardware and software systems. SPORTLOGiQ is targeting hockey, soccer, and basketball. By introducing advanced analytics to 59 professional leagues, the total addressable market becomes $822M. Once the technology is fully developed and capable of being used by individual teams at every level of sport, we will be in a position to serve over 578 million athletes across the three targeted sports. We estimate that this is a $5B market opportunity.

Company HQP benefits: SPORTLOGiQ currently has 10 full time research employees with graduate education. The team will be scaling to 30 employees over the next two years. These are highly educated, highly qualified individuals in a sector where top talent is typically recruited by firms in Silicon Valley. SPORTLOGiQ will use the proposed internships as opportunities to train talented students in Canada for specializing in a highly technical industry.

Canadian labour market benefits: The results of the project will provide jobs in the IT/Software sectors, inject capital into the Canadian economy from the international sport industry, and support the growth of Canada's high tech start-up sector. The proposed research will be conducted with trainees who will emerge as highly qualified personnel with ICT expertise. These students will be in high demand, as evidenced by the PIs' record described above. Most of the proposed budget will support the 13 graduate students and about 9 undergraduate students we will work with over the project. Canada has a shortage of highly skilled ICT workers, which the federal government has aimed to address with extensive HQP funding in the Pan-Canadian AI strategy and the digital superclusters. Graduate research positions in Computer Science and Engineering attract excellent international students to Canada, where they often remain to apply their skills.