Learning Agent Representations for Ice Hockey

Guiliang Liu (**Presenter**), Oliver Schulte, Pascal Poupart,
Mike Rudd, and Mehrsan Javan.

Player Embedding

Embed Player Information

- **Definition:** Modeling the special characteristics of individual players (Learning Agent Representation).

Player Embedding

Embed Player Information

- **Definition:** Modeling the special characteristics of individual players (Learning Agent Representation).
- **Motivation**: e.g., Expected Goal (XG) Prediction.

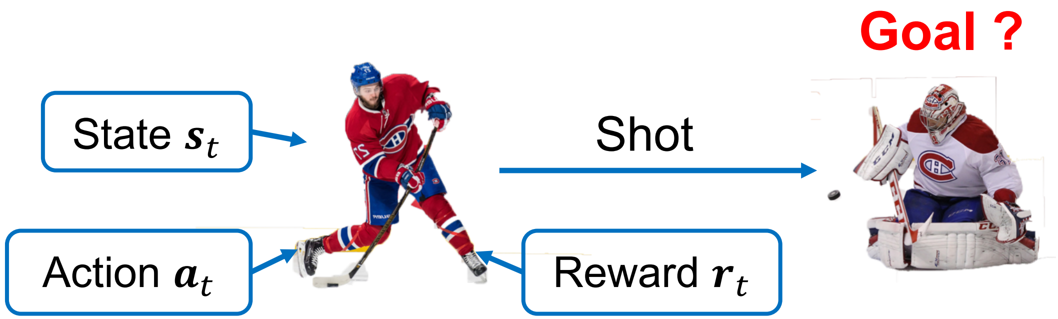

Player Embedding

Embed Player Information

- **Definition:** Modeling the special characteristics of individual players (Learning Agent Representation).
- **Motivation**: e.g., Expected Goal (XG) Prediction.

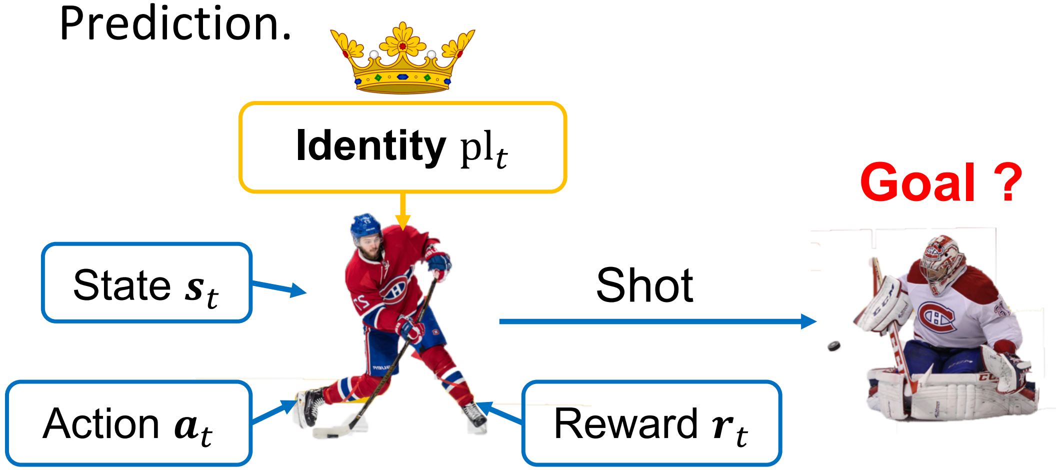

Player Embedding

Embed Player Information

- **Definition:** Modeling the special characteristics of individual players (Learning Agent Representation).
- **Motivation**: e.g., Expected Goal (XG) Prediction.

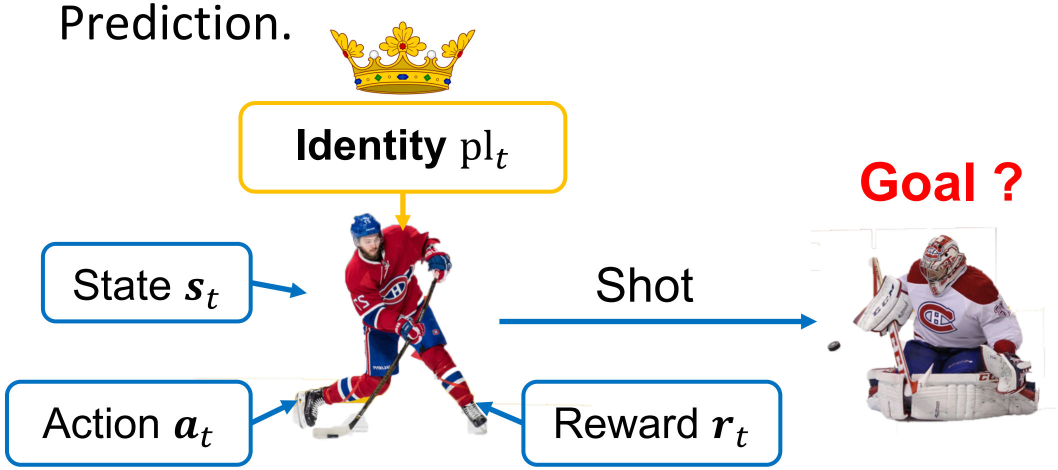

- **Challenges**:
  1) Large agent number (>1k).
  2) Sparse and diverse player observations.

Player Embedding

Embed Player Information

- **Definition:** Modeling the special characteristics of individual players (Learning Agent Representation).
- **Motivation**: e.g., Expected Goal (XG) Prediction.

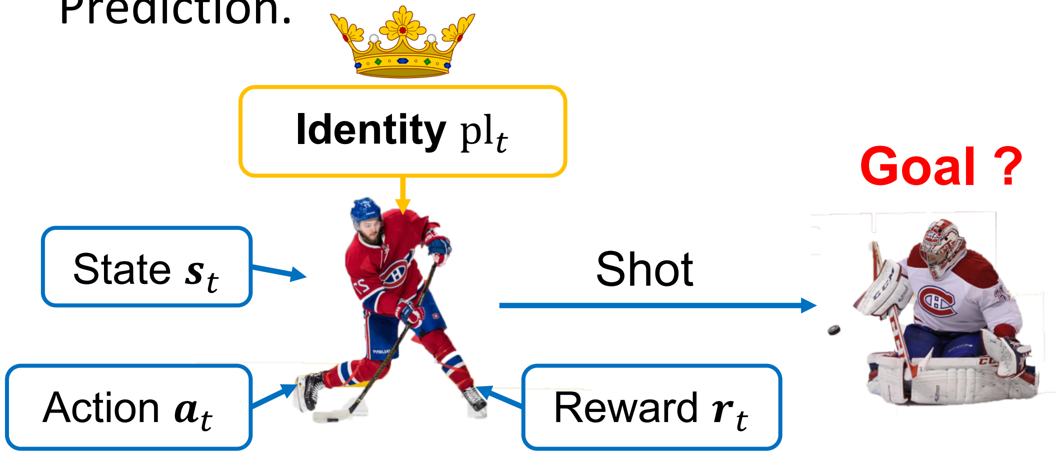

- **Challenges**:
  1) Large agent number (>1k).
  2) Sparse and diverse player observations.

Player Representation via Generation

- **Contextualized representation:** Apply latent variables $\boldsymbol{z}_t$ as a representation of game context:

$$p(pl_t|\mathbf{s}_t, \mathbf{a}_t, \boldsymbol{r}_t) = \int p(pl_t|\mathbf{z}_t)p(\mathbf{z}_t|\mathbf{s}_t, \mathbf{a}_t, \boldsymbol{r}_t)\mathrm{d}\mathbf{z}_t$$

to Learn a *context-aware prior* $p(\boldsymbol{z}_t|\boldsymbol{s}_t, \boldsymbol{a}_t, \boldsymbol{r}_t)$.

Player Embedding

Embed Player Information

- **Definition:** Modeling the special characteristics of individual players (Learning Agent Representation).
- **Motivation**: e.g., Expected Goal (XG) Prediction.

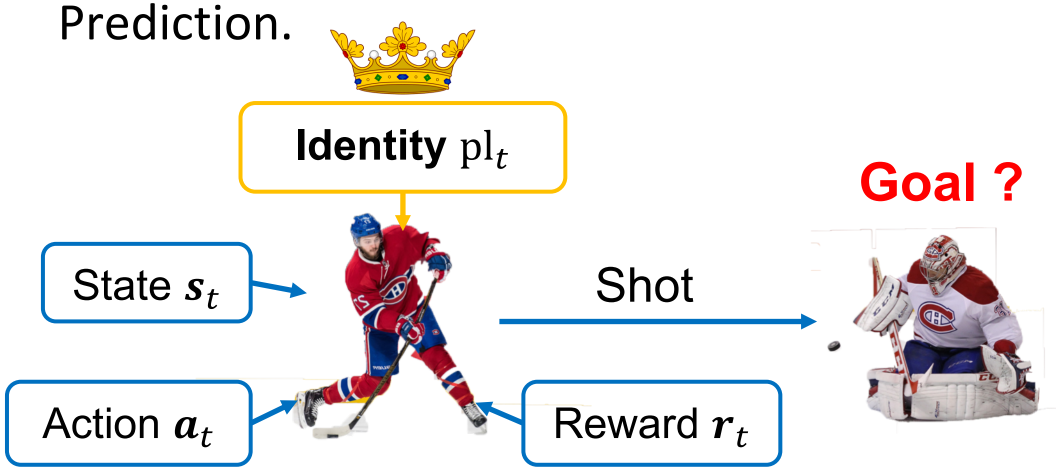

- **Challenges**:
  1) Large agent number (>1k).
  2) Sparse and diverse player observations.

Player Representation via Generation

- **Contextualized representation:** Apply latent variables $\boldsymbol{z}_t$ as a representation of game context:

$$p(pl_t|\mathbf{s}_t, \mathbf{a}_t, \boldsymbol{r}_t) = \int p(pl_t|\mathbf{z}_t)p(\mathbf{z}_t|\mathbf{s}_t, \mathbf{a}_t, \boldsymbol{r}_t)\mathrm{d}\mathbf{z}_t$$

  to Learn a *context-aware prior* $p(\boldsymbol{z}_t|\boldsymbol{s}_t, \boldsymbol{a}_t, \boldsymbol{r}_t)$.

- *After* observing the acting player pl_t, we learn an approximate posterior q$(\boldsymbol{z}_t|\boldsymbol{s}_t, \boldsymbol{a}_t, \boldsymbol{r}_t, pl_t)$ as a *contextualized player representation*:

$$p(pl_t|\mathbf{s}_t, \mathbf{a}_t, \boldsymbol{r}_t) \approx \int p(pl_t|\mathbf{z}_t)q(\mathbf{z}_t|\mathbf{s}_t, \mathbf{a}_t, \boldsymbol{r}_t, pl_t)\mathrm{d}\mathbf{z}_t$$

Model Structure

Variational Recurrent Ladder Agent Encoder (VaRLEA)

- Build a *contextualized* Ladder VAE at RNN cell
  1. Ladder latent variables ($z_{s,t}$, $z_{a,t}$, $z_{r,t}$) for maintaining the *causal dependency* $s \rightarrow a \rightarrow r$.
  2. **Generation**: Estimate *contextualized prior.*
  3. **Inference**: Derive the player representation.

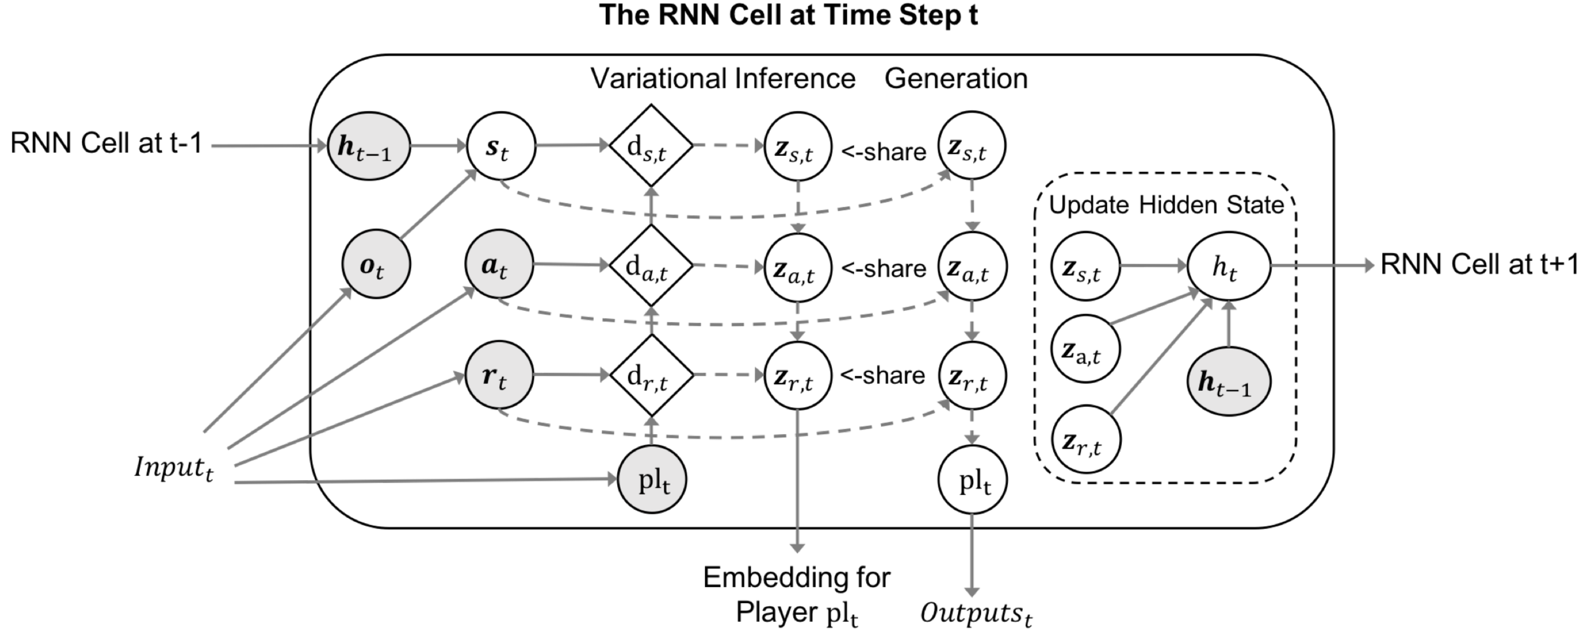

Variational Recurrent Ladder Agent Encoder (VaRLEA)

- Build a *contextualized* Ladder VAE at RNN cell
  1. Ladder latent variables $(z_{s,t}, z_{a,t}, z_{r,t})$ for maintaining the *causal dependency* $s \rightarrow a \rightarrow r$.
  2. **Generation**: Estimate *contextualized prior.*
  3. **Inference**: Derive the player representation.
- **Shrinkage Effect:**
  1. Allow information to transfer between Player observations.
  2. Shrinkage estimator prevents overfitting

The RNN Cell at Time Step t

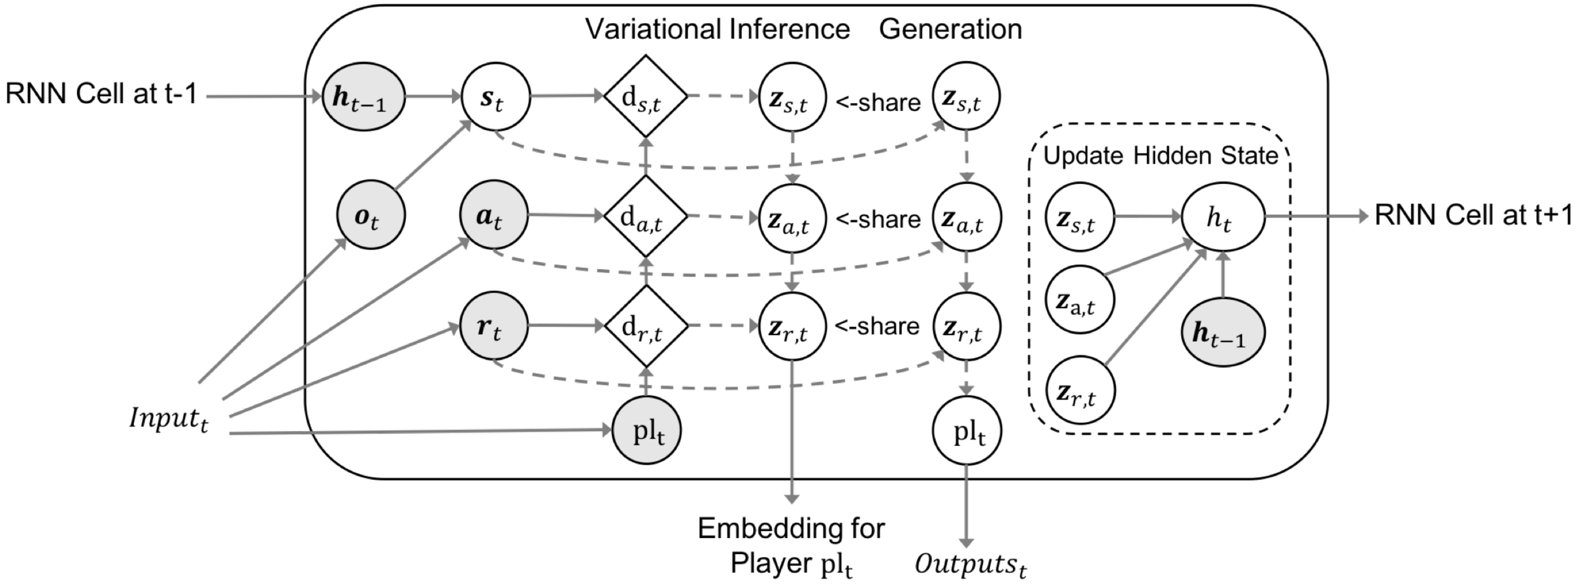

Conclusion

The take-home message:

- The VAE ELBo loss is an effective *shrinkage estimator* for large multi-agent action datasets with
  1. Diversity (many different agents).
  2. Sparsity (some agents with limited observations).
- Modeling player information with player representation improves the performance of downstream applications.