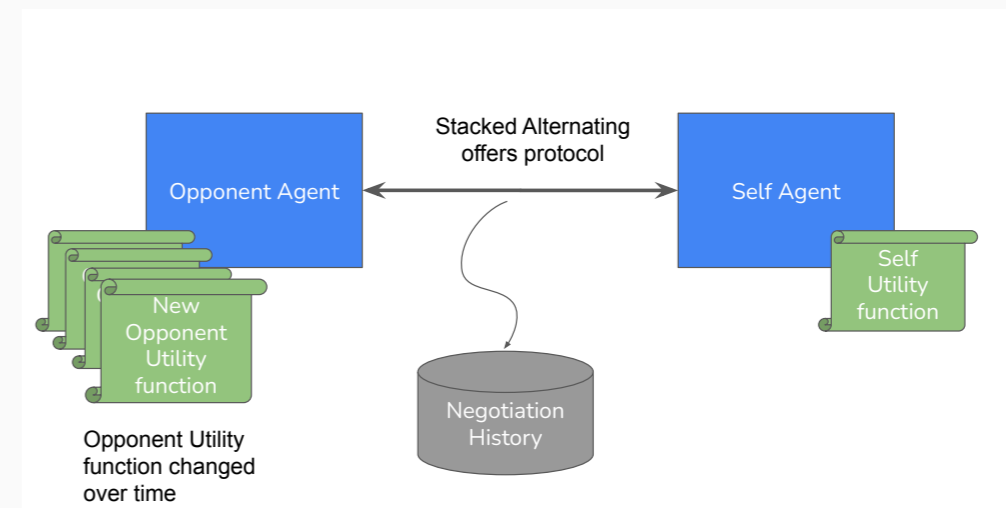


Ayan Sengupta, Shinji Nakadai and Yasser Mohammad

NEC Corporation

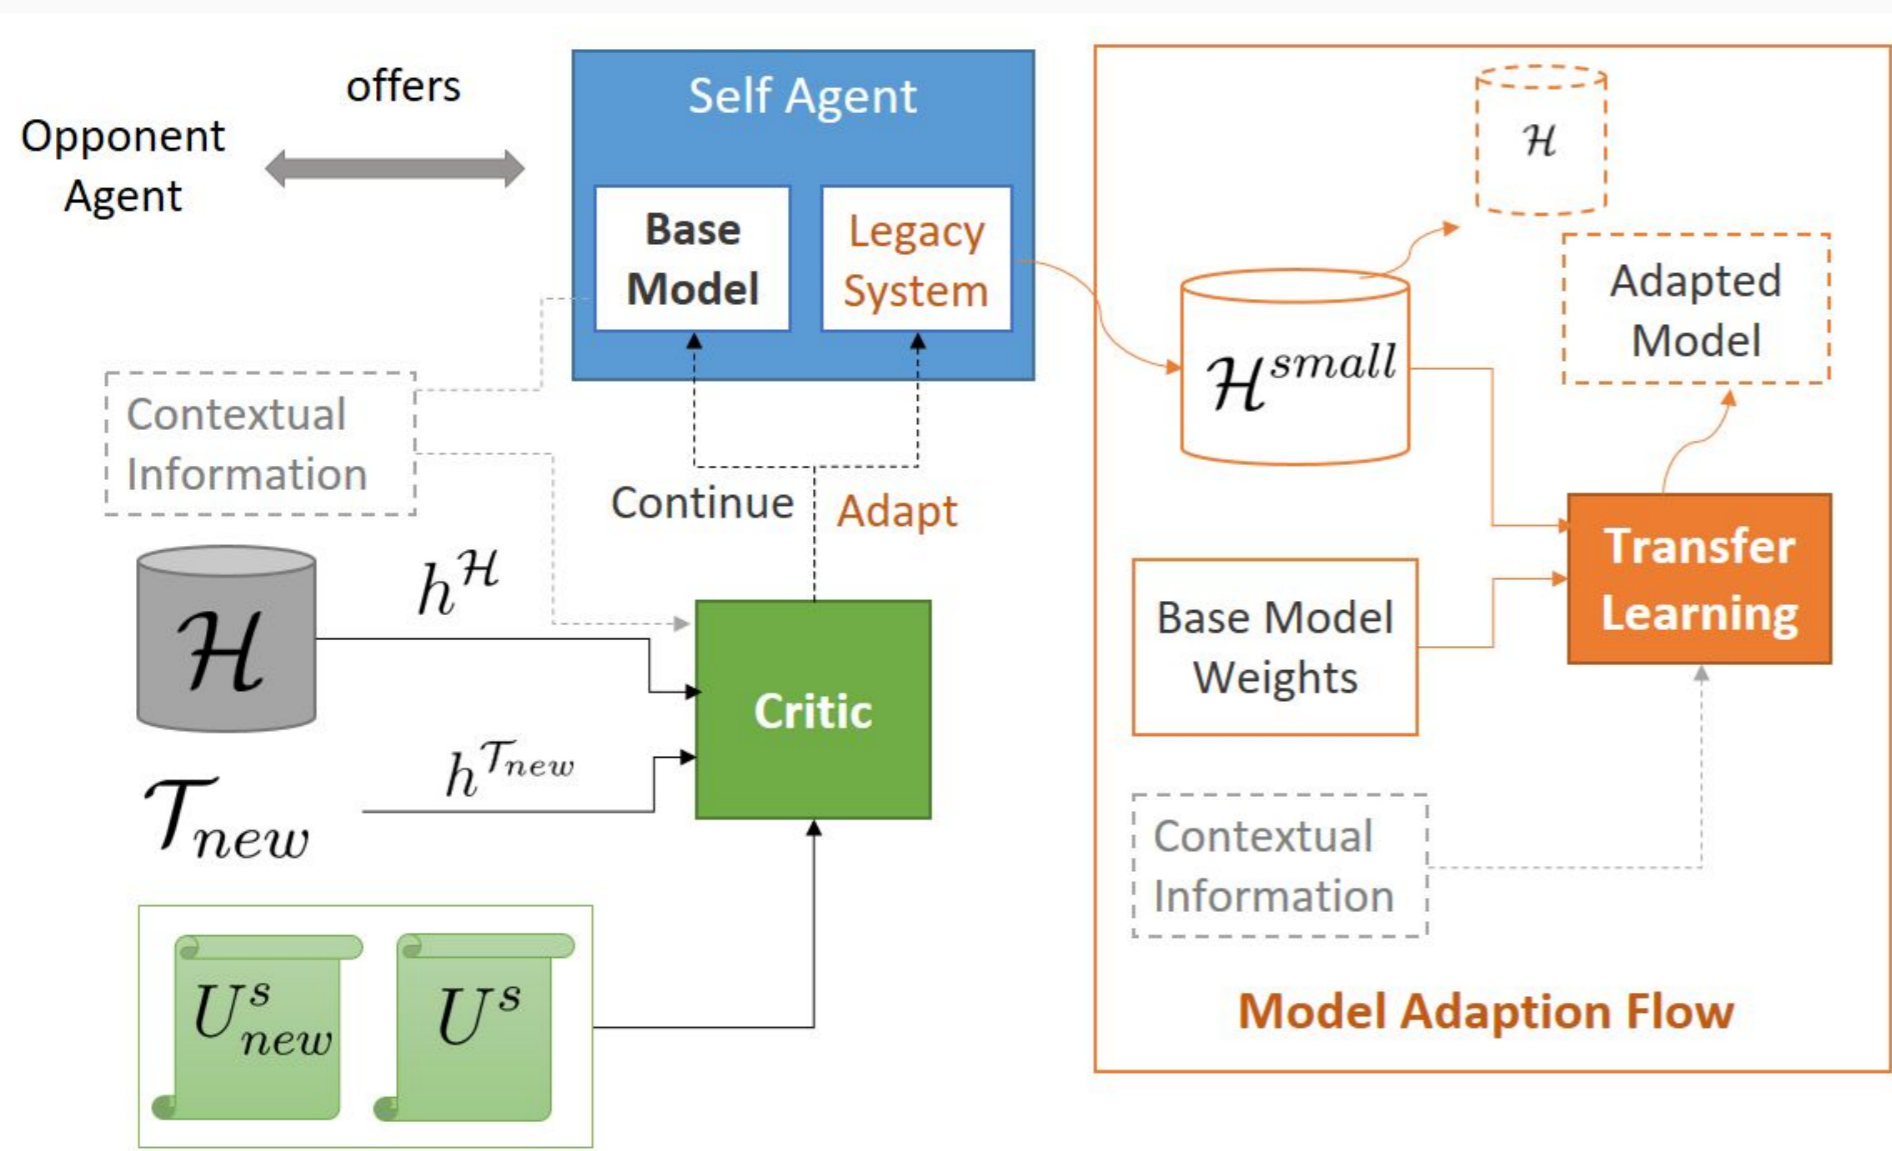
National Institute for Advanced Industrial Science and Technology (AIST)

- Automated negotiation is a ubiquitous method for reaching agreement between industrial partners.
 - In many relevant real-world scenarios, **negotiations repeat with similar but not the same ufuncs.**
- Historical negotiation traces are becoming available due to the increased pace of digitalization specially due to the pandemic.
- We propose a framework for incrementally learning and adapting an end-to-end negotiation policy based on historical negotiation traces.
 - We show that learning is **transferred** to related new situations.
 - The system adapts automatically to **changes in utility functions.**
 - Outperforms** top GENIUS agents on ANAC domains by 6%.



- Takes contextual information into account (e.g. season).
- Trains and adapts with few negotiation traces with no active opponent.
- Your best negotiator in → higher performance out.

Proposed Framework



LegacySystem Top GENIUS Negotiator used for data collection during incremental improvement:

- Atlas3, YX, Ponpoko, NTfT, ParsCat.

BaseModel An LSTM used to learn the negotiation policy.

Critic Decides when to adapt the BaseModel.

Base Model and Transfer Learning

Base Model

Layer (type)	
embedding (Embedding)	
bi_LSTM (Bidirectional)	
dropout (Dropout)	
dense (Dense)	
Total params: 51,040	
Trainable params: 51,040	
Non-trainable params: 0	

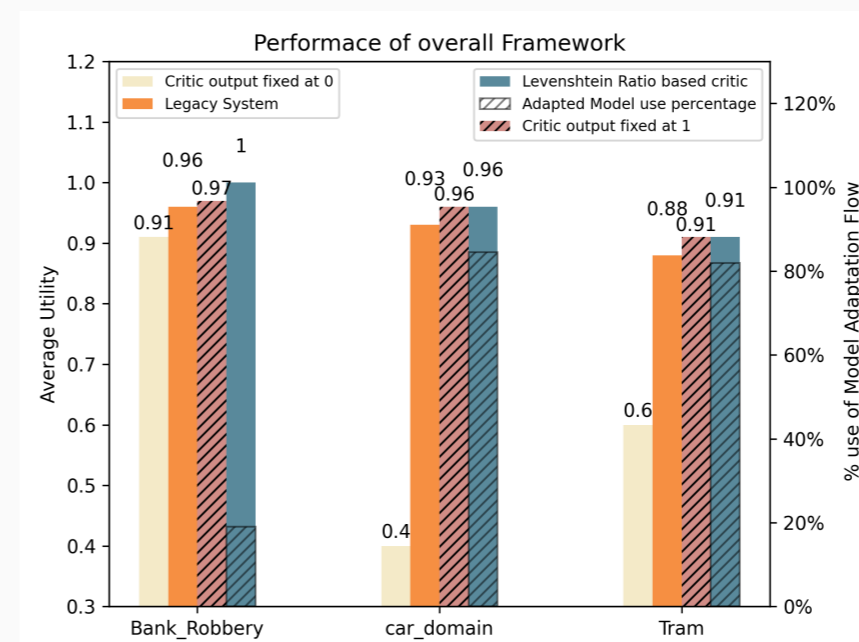
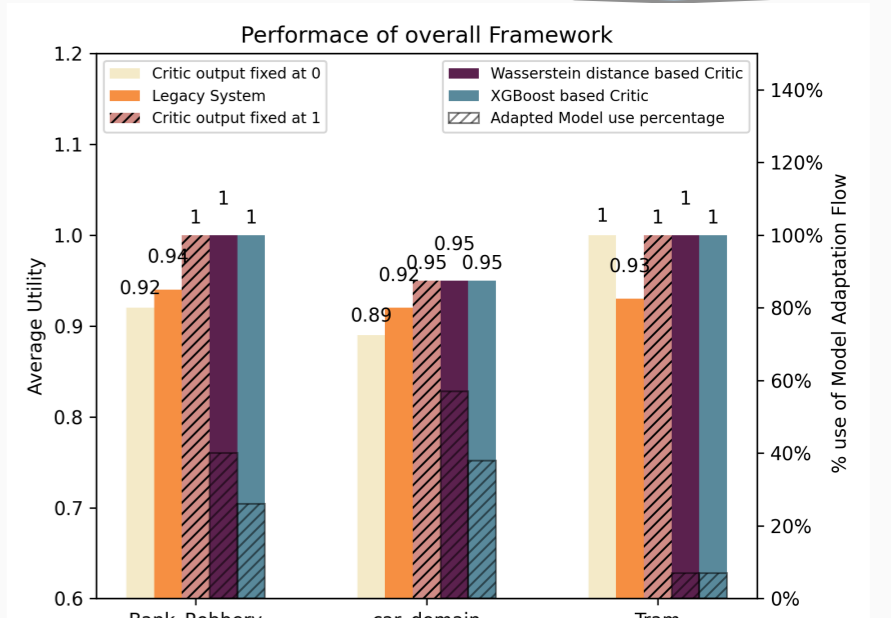
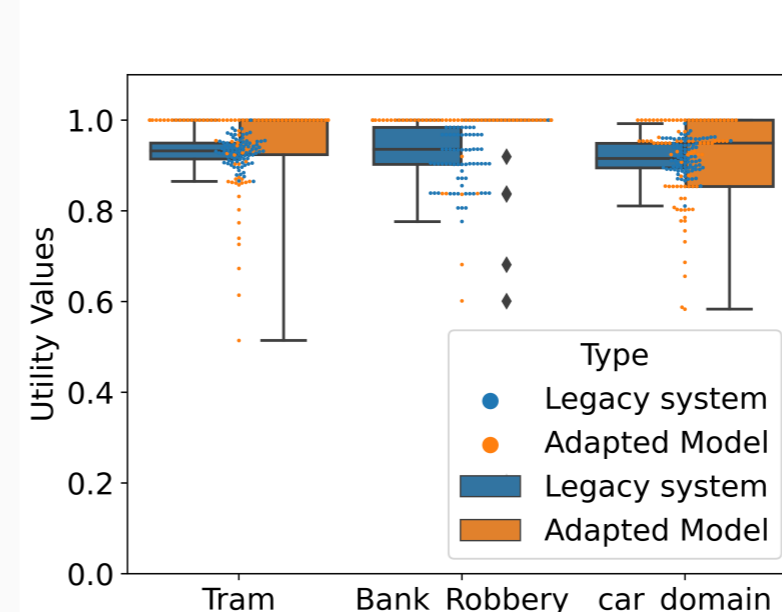
Adapted Model

Layer (type)	
embedding (Embedding)	Retrained
bi_LSTM (Bidirectional)	Shared weights
dropout (Dropout)	
dense (Dense)	Retrained
Total params: 51,040	
Trainable params: 26,640	
Non-trainable params: 24,400	

- The BaseModel is trained using all available traces.
- Input and output encoding layers learn distances in the ufunc space.
- A middle LSTM that *roughly* represents the strategy is not retrained.
- The Critic uses the LegacySystem to collect few (e.g. 10) extra traces when needed and triggers the transfer learning model.

Critic: Adaptation

- A binary classifier (Adapt?).
- Assumes either self or partner ufunc can change.
- Uses the relative frequency of offers in the last few traces to decide whether the partner's ufunc has changed.
- Trained using XGBoost with synthetic data.
 - During cold-start, the Wasserstein distance with an appropriate threshold is also effective.
- Can accommodate changes in self ufunc:
 - Uses Levenshtein distance to measure the change between the new and old ufuncs.



- Performance 6% ↑.
- Ablation: Critic is effective.
- XGBoost, Wasserstein distance are both effective but with different retraining rates.