Statement of Research Interests

Jaesik Choi

Intelligent systems are enhancing many aspects of daily life. Accurate environmental forecasting with large-scale sensor networks [Choi *et al.*, 2011b; Choi and Amir, 2012; Xu *et al.*, 2012], dependable protection of safety-critical systems [Choi *et al.*, 2010b; 2011c; Hajishirzi *et al.*, 2009; Yoon *et al.*, 2012], content-based multimedia retrieval [Choi *et al.*, 2008; 2012] and controlling robot arms for human interface devices [Choi and Amir, 2007; 2009] are only small samples of possible applications where intelligent systems may change the physical world. Discrete and continuous elements interact closely with others in the real-world systems.

Some of the key research challenges involve learning models from physical systems and inference with the models by calculating various interesting probabilities. In general, it is very difficult to learn and inference with large-scale physical systems. Nonetheless, my research has shown that efficient state estimations over large-scale hybrid continuous-discrete systems are within reach (e.g., [Choi *et al.*, 2010a; 2011a; 2011b; Choi and Amir, 2012; Choi, 2012]).

In the following, I will describe the results of the research I have done during my thesis study, the projects that I plan to develop in the future, and the research principles that I hold.

1 My Thesis Work

My doctoral work at University of Illinois at Urbana-Champaign [Choi, 2012] introduces novel ways of learning and inference for hybrid models with natural language-like knowledge representations and continuous (possibly noisy) inputs from physical systems. The new methods are applied to estimate the changes in various physical systems from forecasting environmental changes [Choi and Amir, 2012; Xu *et al.*, 2012] to protecting safety-critical systems from intrusions [Yoon *et al.*, 2012].

The main contributions of my thesis are algorithms that maintain compact representations during inference and filtering procedures over large numbers of random variables. I showed that two random variables have the same variances and covariances after filtering steps, when the two variables have the same relationships with others in a linear dynamic system [Choi *et al.*, 2010a; 2011b]. This fact holds even when the two random variables have different observations. In this way, my algorithm enables the preservation of compact structures over clustered random variables as shown in Figure 1. This new principle paves the way to build Relational Kalman Filtering (RKF) which enables scaling the exact vanilla Kalman filter from thousands of variables to billions.

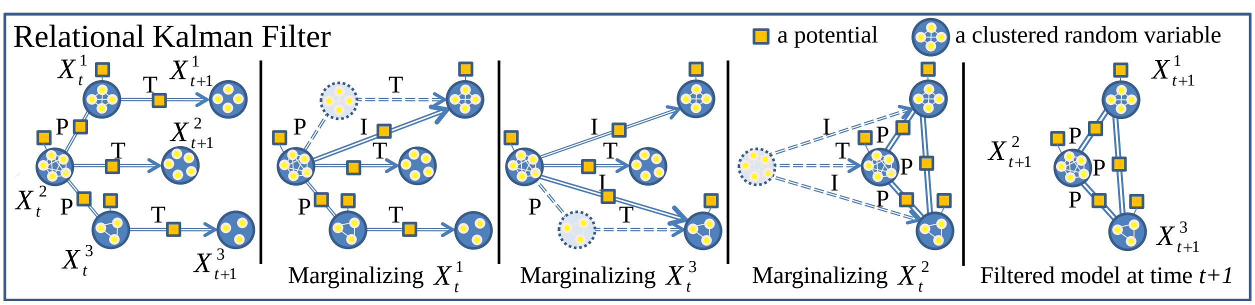

Figure 1: The relational representation dramatically eliminates the need for redundant potentials, given three sets of variables X_t^1, X_t^2 and X_t^3. Hence, representation and filtering can be scalable to large models. Note that the conventional (propositional) representations and the vanilla Kalman filter are not suitable for large-scale dynamic systems.

Also, I took further steps towards fundamental understanding of large-scale hybrid models. I showed that inference with commonly used natural language-like knowledge representations (or relational models) for large populations, can be accurately approximated by simpler representations such as Gaussians [Choi *et al.*, 2011a] and variational models [Choi and Amir, 2012]. These foundings led ways to build scalable inference algorithms for large-scale relational models as shown in Figure 2.

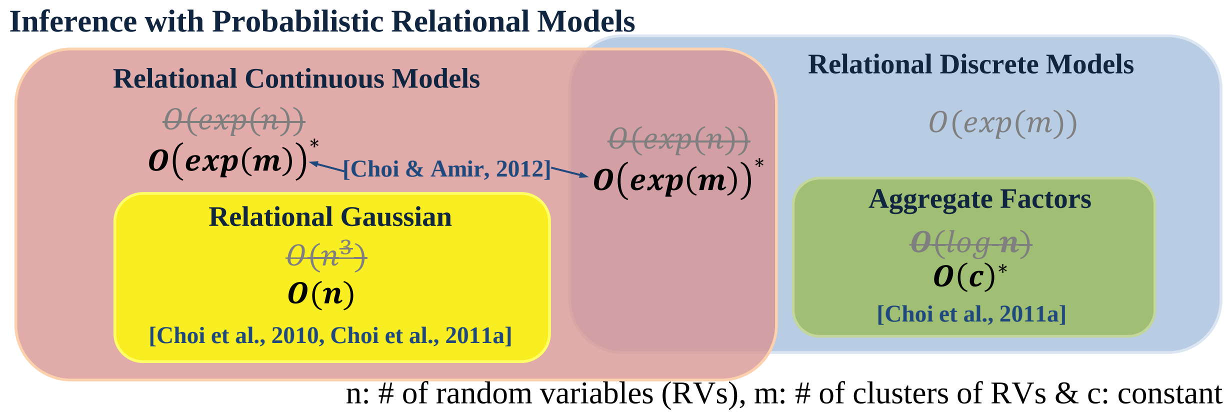

Figure 2: This diagram overviews the theoretical contributions of my thesis work in terms of computational complexities. Complexities in bold represent new results. '$*$' refers to an approximation.

2 Future Plans

I plan to address research challenges, which seamlessly unify natural language-like knowledge representations and physical systems. Previously, research in traditional Artificial Intelligence (AI), Logic, and formal verification focuses sorely on combinatorial structures with discrete variables in knowledge representations. Meanwhile, research in Machine Learning (ML) and Robotics focuses more on continuous distributions or processes in physical systems. However, many real-world applications involve both aspects. My future research will focus on developing new learning and inference algorithms for such hybrid models.

2.1 Inference with Large-Scale Models

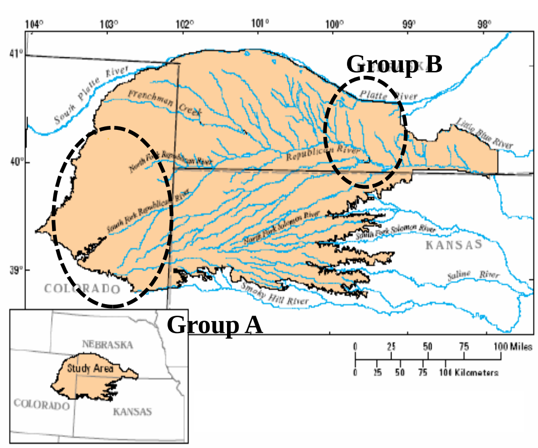

Figure 3: The Republican River Basin.

Large-scale models such as monitoring climate data with sensors over a large area, bring several challenges for existing inference algorithms. Learning and inference problems with such large, complex systems are intractable for existing methods. Relational models compactly represent relationships among a large number of random variables. I will investigate further on my expertise – efficient inference algorithms over relational hybrid models (large-scale graphical models) and machine leanring algorithms for real-world environment data [Choi *et al.*, 2010a; 2011b; Choi and Amir, 2012].

The main idea of the relational models is to cluster random variables for efficient, accurate estimations. For example, it has been shown that clustering water wells, e.g., Group A and B, improved the accuracy of groundwater models in the Republican River Basin [Xu *et al.*, 2012].

I will further strive to discover novel ways of exact and approximate inference over the relational models of physical systems. One of such examples is the *Relational Rao-Black Particle filter* which is an approximate, scalable filter for non-linear dynamic systems. I will bring the principles of my previous work [Choi *et al.*, 2011b] when designing such new systems.

2.2 Spatio-Temporal Analysis with Exchangeable Random Variables

One of salient characteristics of probabilistic relational models, e.g., First-Order Probabilistic Models and Markov Logic Networks (MLNs), is that groups of random variables tend to be exchangeable[1] [Choi and Amir, 2012]. For discrete variables, such exchangeable variable can be represented by histograms.

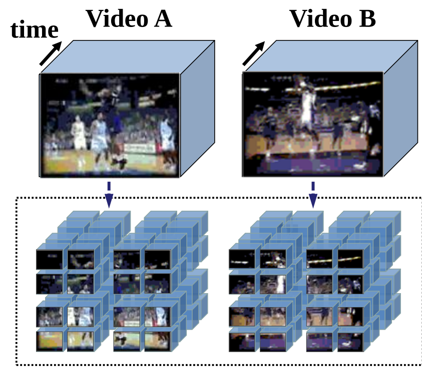

Figure 4: An illustration of the Spatial-Temporal Pyramid Matching (STPM) kernel.

Representing features in histograms is already widely used in computer vision when measuring similarities among multimedia data (images and videos). My work, Spatio-Temporal Pyramid Matching (STPM) [Choi *et al.*, 2008]), reveals a new way of building histograms for video data when spatial features and temporal features interact tightly. The STPM and its variants are currently one of state-of-the-art methods in recognizing human actions out of videos [Choi *et al.*, 2012].

This research can be a decent starting point to discover new insight on how to design a visual recognition framework using language-like knowledge representations. I will focus on exploiting human knowledge in forms of probabilistic first-order representations using such exchangeable random variables. I believe that rich combinatorial structures of the representations can help to recognize visual objects.

2.3 Combining Planning and Motion Planning to Manipulate Smartdevices

Another line of research that I hope to pursue is combining motion planning and AI planning for navigating and manipulating in a human-oriented environment [Choi and Amir, 2007; 2009]. For example, smartphones normally include human-friendly interfaces. However, manipulating such devices is challenging for robots because assumptions behind motion planning and AI planning are fragile for such manipulation tasks. C-space (configuration space) is not expressive enough to represent the control logistics of smartphones. In addition, AI planning assumes a set of well-defined actions (e.g. touching a button). However, such actions are easily affected or changed using many physical factors, such as the configuration of a device and the constraints of a robot arm.

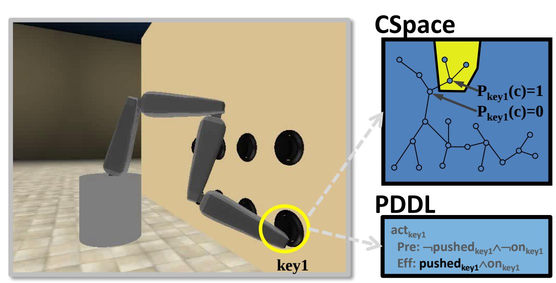

Figure 5: An example of combining an AI planner (e.g., PDDL) and a motion planner (e.g., C-Space)

I will address this issue using my work [Choi and Amir, 2009] and focus on learning general-purpose actions from a motion planning algorithm that takes kinematic constraints into account. The motion planning will be used to verify feasible actions on the user-specified actions. These algorithms will make it possible to manipulate devices designed for human, such as iPhones, with robotic arms.

3 Conclusions

The driving force behind my research is to build human-level intelligent systems. I believe that this goal is best pursued by efforts to understand principles of learning and reasoning. Thus, it is crucial to collaborate

[1]Random variables are exchangeable, when a joint probability of the variables is invarient to any permutation of variables.

with multidisciplinary research areas, e.g., probabilistic reasoning, machine learning, cognitive science, motion planning, and sensing. I will keep moving forward to understand computational and modeling principles of algorithms, and apply the results to improve various real-world physical systems.

References

[Choi and Amir, 2007] Jaesik Choi and Eyal Amir. Factor-guided motion planning for a robot arm. In *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2007)*, pages 27–32, 2007.

[Choi and Amir, 2009] Jaesik Choi and Eyal Amir. Combining planning and motion planning. In *2009 IEEE International Conference on Robotics and Automation (ICRA 2009)*, pages 238–244, 2009.

[Choi and Amir, 2012] Jaesik Choi and Eyal Amir. Lifted relational variational inference. In *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence (UAI 2012)*, 2012.

[Choi *et al.*, 2008] Jaesik Choi, Won J. Jeon, and Sang-Chul Lee. Spatio-temporal pyramid matching for sports videos. In *Multimedia Information Retrieval (ACMMIR 2008)*, pages 291–297, 2008.

[Choi *et al.*, 2010a] Jaesik Choi, Eyal Amir, and David Hill. Lifted inference for relational continuous models. In *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence (UAI 2010)*, pages 126–134, 2010.

[Choi *et al.*, 2010b] Jaesik Choi, Jay Pujara, Vishwanth Tumkur Ramarao, and Ke Wei. Identifying ip addresses for spammers. Patent. US 7849146, 2010.

[Choi *et al.*, 2011a] Jaesik Choi, Rodrigo de Salvo Braz, and Hung H. Bui. Efficient methods for lifted inference with aggregate factors. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2011)*, pages 1030–1036, 2011.

[Choi *et al.*, 2011b] Jaesik Choi, Abner Guzman-Rivera, and Eyal Amir. Lifted relational kalman filtering. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI 2011)*, pages 2092–2099, 2011.

[Choi *et al.*, 2011c] Jaesik Choi, Ke Wei, and Vishwanath Tumkur Ramarao. Filter for blocking image-based spam. Patent. US 8055078, 2011.

[Choi *et al.*, 2012] Jaesik Choi, Ziyu Wang, Sang-Chul Lee, and Won J. Jeon. A spatio-temporal pyramid matching for video retrieval. under review in Computer Vision and Image Understanding (CVIU), 2012.

[Choi, 2012] Jaesik Choi. *Lifted Inference for Relational Hybrid Models*. Phd thesis, University of Illinois Urbana-Champaign, 2012.

[Hajishirzi *et al.*, 2009] Hannaneh Hajishirzi, Afsaneh Shirazi, Jaesik Choi, and Eyal Amir. Greedy algorithms for sequential sensing decisions. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI 2009)*, pages 1908–1915, 2009.

[Xu *et al.*, 2012] Tiangfang Xu, Albert J. Valocchi, Jaesik Choi, and Eyal Amir. Improving groundwater flow model prediction using complementary data-driven models. In *XIX International Conference on Computational Methods in Water Resources (CMWR 2012)*, 2012.

[Yoon *et al.*, 2012] Man-Ki Yoon, Sibin Mohan, Jaesik Choi, Jung-Eun Kim, and Lui Sha. Securecore: A multicore based intrusion detection architecture for real-time embedded systems. submitted – under review, 2012.