# Deep Amortized Inference for Probabilistic Programs using Adversarial Compilation

Mahdi Azarafrooz
m.azarafrooz@gmail.com

## Summary

We propose an amortized inference strategy for probabilistic programs, one that learns from past inferences to speed up the future inferences. Our proposed inference strategy is to train neural guidance programs via a minimax game, with the probabilistic program as a *correlation device*. From a game-theoretical vantage point, the role of a correlation device is to enforce better outcomes by sharing information between players. The shared information, in our case, is the execution trace, which gets used for computation of payoffs in the minimax game.

## Amortized Inference

PPL expresses the generative model as:

$$p(\mathbf{y},\mathbf{x}) = p(\mathbf{y}|\mathbf{x}) \prod_i p(x_i|\mathbf{x}_{<i}) \quad (1)$$

The amortized inference strategy in [2] is to learn a neural guidance program $q$ upfront so that at inference time, sampling from $q$ is both fast and accurate. It is an approximation of $p(\mathbf{x}|\mathbf{y})$ using side neural computations:

$$q(\mathbf{x}|\mathbf{y};\phi) = \prod_i q(x_i; D_i(\mathbf{y}, \mathbf{x}_{<i};\phi)) \quad (2)$$

where $D_i$ is a neural network.

## Conclusion

The literature of generative models are full of diverse and sometimes even opposing approaches such as sampling, sequential decision making, searching in minimax game, variational optimization, and transformational compilation of probabilistic programs. Our proposal brings all these approaches in one ecosystem where not only all these different views coexist but also enhance each other.

## References

[1] Wingate, D. & Stuhlmüller, A and Goodman, N.D . (2011) Lightweight Implementations of Probabilistic Programming Languages Via Transformational Compilation, *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, AISTATS*
[2] D. Ritchie, P. Horsfall, Noah D Goodman (2016) Deep Amortized Inference for Probabilistic Programs *arXiv preprint arXiv:1610.05735*
[3] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, Yoshua Bengio. (2017) Generative Adversarial Nets *Advances in Neural Information Processing Systems 27*
[4] Robert J. AUMANN (1974) subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics.*
[5] L. Yu, W. Zhang, J. Wang, Y. Yu (2017) SeqGAN: Sequence Generative Adversarial Nets with Policy Gradient *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17)*
[6] S. Hochreiter and J. Schmidhuber. (1997). Long Short-Term Memory. *Neural Comput. 9.*
[7] K. Cho, B. v. Merrienboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. (2014) Learning Phrase Representations using RNN EncoderDecoder for Statistical Machine Translation.

## Neural Game Theoretic Preliminaries

**Generative Adversarial Networks [3]** Considers 2 neural networks $G_\theta$ (generator) and $D_\phi$ (adversary) as the players of a 2 player minimax game with Jenson-Shannon distance as utility.

**Nash Equilibrium** : A joint action strategy $\mathcal{A} = (\theta, \phi)$ is called Nash equilibrium (NE) if no neural network achieves a better loss by deviating from the equilibrium unilaterally.

**Correlated Nash Equilibrium** : Let $(\Omega, \pi)$ be a countable probability space. For each player $i$, let $\mathcal{P}_i$ be his information partition, $\rho_i$ be $i$'s posterior and let $s_i : \Omega \to A_i$ be the non-linear neural operators, assigning the same value to states in the same cell of $i$'s information partition $\mathcal{P}_i$. Then $((\Omega,\pi), \mathcal{P}_i, s_i)$ is a correlated equilibrium CE if for every player $i$ and for every strategy modification $\psi_i$:

$$\sum_{w\in\Omega} \rho_i(w) u_i(s_i(w), s_{-i}(w)) \geq \sum_{w\in\Omega} \rho_i(w) u_i(\psi_i(s_i(w)), s_{-i}(w)) \quad (3)$$

In other words, $((\Omega,\pi), \mathcal{P}_i, s_i)$ is a CE if not player can improve his/her own expected utility via a strategy modification.

A NE is a CE such that joint posterior $\rho$ is a product distribution; that is $\rho = \prod_{i=1}^{n} \rho_i$.

Since $\rho(w)$ is a probability distribution , $\rho(w) \geq 0$ and $\sum \rho(w) = 1$ computing a correlated equilibrium only requires solving a linear program whereas solving a Nash equilibrium requires finding its fixed point completely.

## Adversarial Amortized Inference

The adversary 's goal is to maximize its payoff:

$$V_{D_\phi} = \underset{\mathbf{x}\sim \underbrace{\boxed{\mathcal{F}(\mathbf{y})}}_{\text{PPL as correlation device}}}{\mathbb{E}} [\log D(\mathbf{y},\mathbf{x};\phi)] + \underset{\mathbf{x}\sim G(\mathbf{y};\theta)}{\mathbb{E}}[\log(1 - D(\mathbf{y},\mathbf{x};\phi))]$$

Generator's payoff is defined as expected end payoff formulated in Eq. 4, similar to [5].

$$V_{G_\theta} = -\sum_i G(\mathbf{y}, \mathbf{x}_{<i};\theta) Q_{D_\phi}^{\mathcal{F}^{G_\theta}}(s = \mathbf{x}_{<i}, a = x_i) \quad (4)$$

where $Q_{D_\phi}^{\mathcal{F}^{G_\theta}}(\mathbf{x}_{<i}, a)$ is the reward function for taking compilation action $a$, and then following the inferred execution trace $\mathcal{F}(\mathbf{y})$ with $G_\theta$ as its proposal distribution:

$$Q_{D_\phi}^{\mathcal{F}^{G_\theta}}(s = \mathbf{x}_{<i}, a = x_i) = \quad (5)$$

$$\begin{cases} D(\mathbf{y},\mathbf{x};\phi), \mathbf{x} \sim \mathcal{F}^{G_\theta} & if\ i < N \\ D(\mathbf{y},\mathbf{x};\phi) & else \end{cases}$$

**Proposition** 1 [4] *There may exist mutually advantageous equilibrium points for a 2-person minimax game if we permit both correlation and subjectivity.*

A theoretical comparison between amortized inference strategy of [2] and ours is not easy. This is because [2] minimizes the KL distance, whereas we formulate a minimax search for minimum JD distance. However, we would like to reiterate that, unlike [2], our proposed amortized inference strategy is not limited by fully-factored representational formats.

We are currently integrating this inference strategy with the Anglican PPL framework.