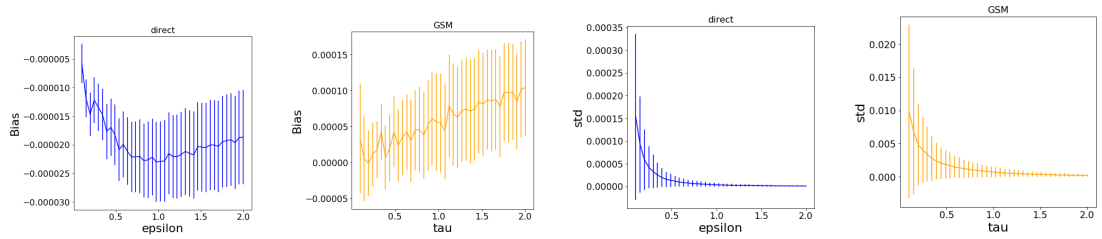


Figure 1: Bias-Variance tradeoff

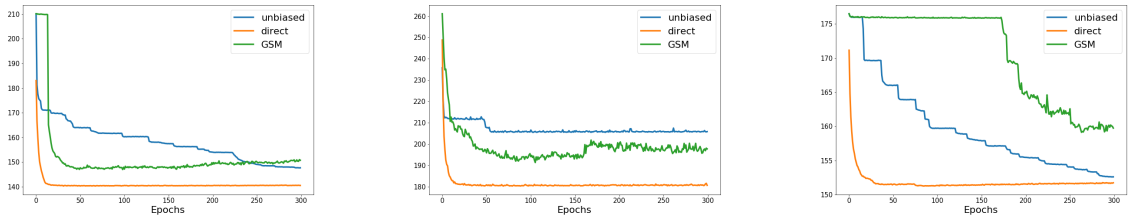


Figure 2: Test loss for $K = 50$ (right: MNIST, middle: Fashion-MNIST, left: Omniglot

We thank the reviewers for their time and appreciate the feedback. Are happy all reviewers agree the work is significant and novel ("Novel approach... High significance. General problem, area of active research"; "seems significant and original"; "This contribution is important").

**[R1]** We thank you for your notes and we will use the extra page to provide a more thorough and clear summary of our components, especially direct loss minimization. We will also make a clear distinction between the model parameters and the distributions and we will consolidate our terms about Gumbel-Max trick and Gumbel-Max perturbation models, thank you for emphasizing it in your review.

**[R3]** We thank you for your suggestions and we will revise the work accordingly as well as include an algorithm of the method in the paper. We will also add more results for the CelebA in the supplementary material — we did not include them in this response due to lack of space.

**[R2] $\epsilon$ and $\tau$ are different variables in quite different settings:** We agree they differ but since their order of magnitude is different, we lose information when placing them in the same plot. In Figure 1 we separated all four options and added error bars. **In the direct loss paper, there are no gradients inside the expectation:** Their work considers linear models $\phi(x, y; w) \stackrel{def}{=} \langle w, \hat{\phi}(x, y) \rangle$ (note, $\phi(x, y; w)$ in our notation is $s_w(x, y)$ in their notation in their paragraph 2) while our work considers non-linear models. Therefore, in their setting $\nabla_w \phi(x, y; w) = \hat{\phi}(x, y)$ and in our setting we consider the general gradient notation $\nabla_w \phi(x, y; w)$. **Puzzled by the final step of your proof:** Following the complete proof in the appendix: $\partial_w G(w, \epsilon) = \mathbb{E}_\gamma[\nabla_w \phi(x, z^{\epsilon\theta+\phi+\gamma}; w)]$. Since we show that $G(w, \epsilon)$ is smooth, then $\partial_w G(w, \epsilon)$ is smooth and its derivative with respect to $\epsilon$ existsThis derivative is defined by the limit operation. **Puzzled by the fact that your unbiased estimator is beaten by direct and even GSM as k goes up:** We agree that this result is surprising and we also provide code so this behavior can be verified. In Figure 2 in this response we present plots for $k = 50$ to complement the results in the paper ($k = 10$). We see that: (i) the gap is the largest for Fashion-MNIST. (ii) when we let the unbiased algorithm to run for long enough the gap between the unbiased minimization and the GSM and direct is getting smaller (iii) the unbiased algorithm sometimes reduces the test loss in steps, so it is quite possible that if we let the unbiased to run for longer we might see eventually more reductions in the test loss till its performance is the same as direct and GSM. (iv) It is certainly possible that this gap will remain even if we run the algorithms for longer time, as these algorithms optimize a non-convex function. We conjecture that if this is the case then it might be on par with the difference between gradient descent (unbiased) and stochastic gradient descent (direct, GSM) that more strongly pursue direction of decent. We will include all plots ($k = 10, 20, .., 50$) to make this behavior apparent. **Loss for direct and GSM are higher for k=50 than k=40 for Omniglot:** Perhaps with the added parameters a worse minimum was found. We kindly point out that the direct and GSM still improve in $k = 50$ over $k = 40$ in MNIST and Fashion-MNIST. **I would love to have some comparisons for Figure 5 and 4 to your baselines:** These figures are qualitative and generally suggest that supervision helps and provide motivational example for the structured setting. Quantitive results appears in Table 2 and Figure 3. **I'd appreciate standard error and averaging over multiple random seeds:** We ran this experiment and omitted due to lack of space. We will add it to the supplementary material. We will follow the suggestions about the clarity.