Inference & Introspection in DLGMs

Rahul G. Krishnan[1], Matthew Hoffman[2]

[1] New York University [2] Adobe Research

Introduction

- Deep Latent Gaussian Models (DLGMs) have seen much work on domains with dense features, less so with sparse features,
- **Inference:** Two methods to improve inference (and subsequently, learning) in these models,
- **Introspection of the Learned Latent Space:** Develop a method to directly visualize how much of the latent space is used
- **Introspection for Representation Learning:** Our method naturally yields feature vectors whose properties we study

Inference & Introspection

Inference Networks

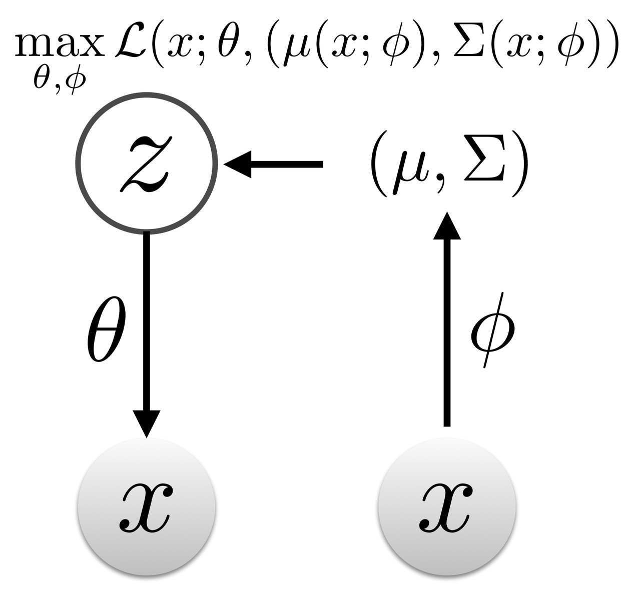

Optimizing local $(\mu(x), \Sigma(x))$

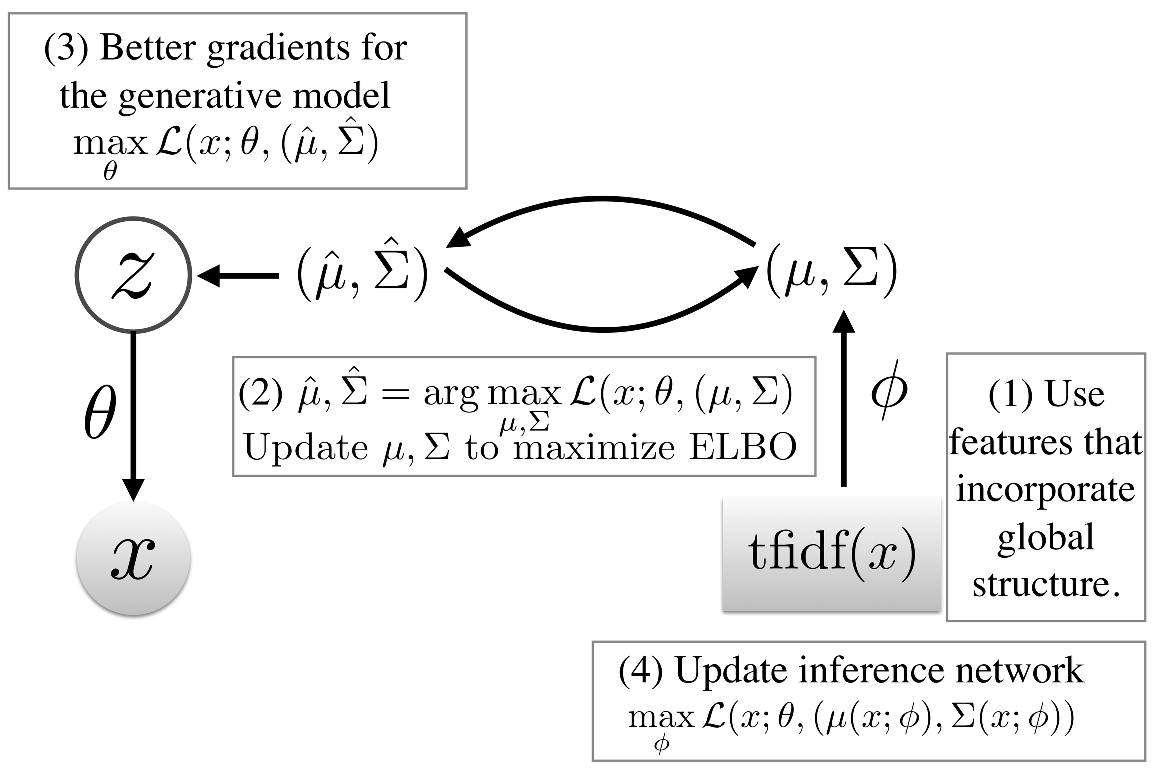

Jacobian Vectors

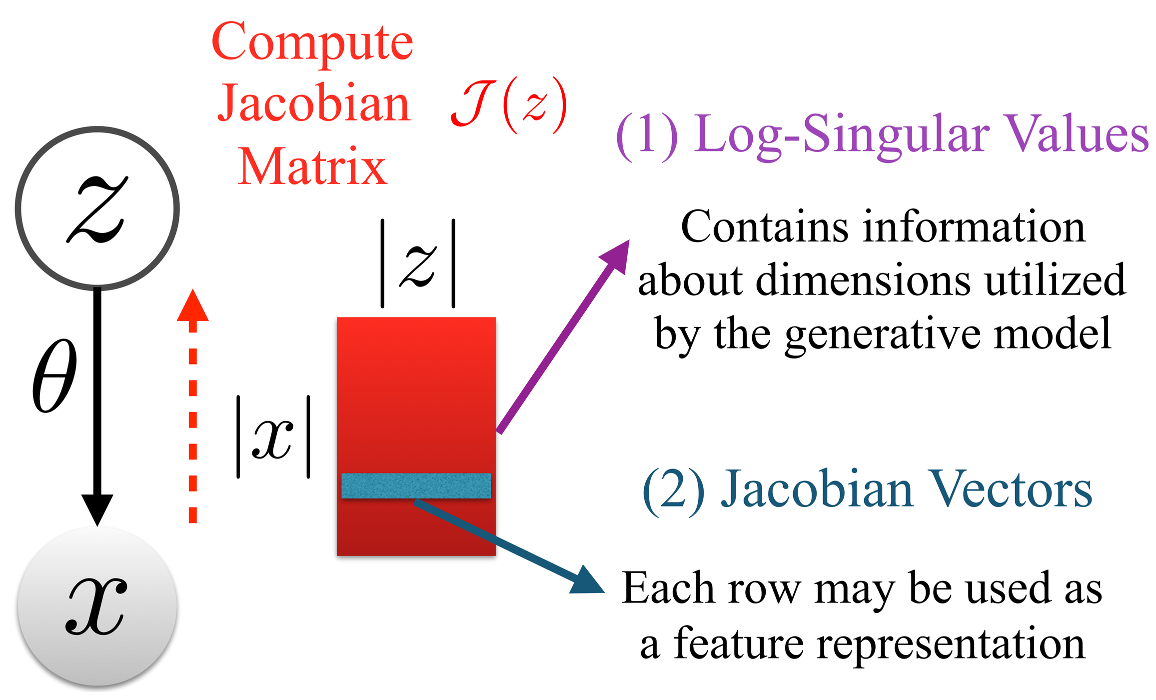

Experimental Results

Model	$\lvert z\rvert$	20News	RCV1-v2
LDA	200	1058	1142
SBN	50	909	784
NVDM	50	836	563
NVDM	200	852	550

Table: **Test Perplexity:** (Above) Baselines, (Below) DLGMs on text data with $|z| = 100$. $M = 1, 100$ (additional steps optimization of $(\mu(x), \Sigma(x))$)

DLGM	20News		RCV1-v2	
	M1	M100	M1	M100
1-M1-norm	964	816	498	479
1-M100-norm	1182	831	485	453
3-M1-norm	1040	866	408	360
3-M100-norm	1341	894	378	329
1-M1-tfidf	895	**785**	475	453
1-M100-tfidf	917	792	480	451
3-M1-tfidf	1027	852	391	346
3-M100-tfidf	1029	833	377	**327**

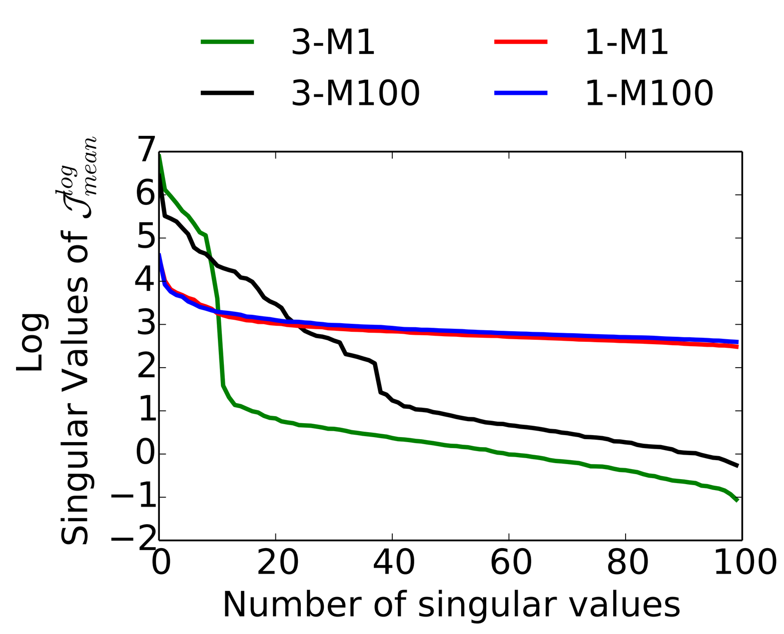

Figure: RCV2

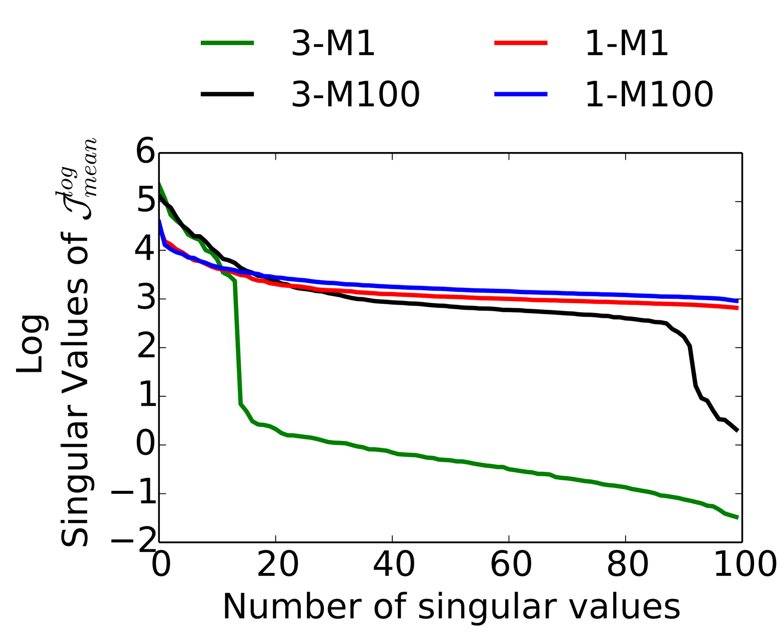

Figure: Wikipedia

Table: **Word Embeddings (Nearest Neighbors):** We visualize rows of $\mathcal{J}(z)$ and show that they preserve semantic structure. Here we show the nearest neighbors of word embeddings (excluding plurals).

Query	Neighborhood
intelligence	espionage, secrecy, interrogation
zen	dharma, buddhism, buddhas, meditation
artificial	artificially, molecules, synthetic, soluble
military	civilian, armys, commanders, infantry

Query	Neighborhood
Superman II	Superman: The Movie, Superman III, Superman IV: The Quest for Peace
Casablanca	Citizen Kane, The Treasure of the Sierra Madre, Working with Orson Welles, The Millionairess
Princess Bride	The Breakfast Club, Sixteen Candles, Groundhog Day, Beetlejuice
12 Angry Men	To Kill a Mockingbird, Rear Window, Mr. Smith Goes to Washington, Inherit the Wind

Table: **Movie Neighbors:** On a model trained on the Netflix dataset, using the rows of the Jacobian as features for movies, we visualize some of the closest neighbors found to movies whose title is displayed on the column on the left

Conclusion

- **Incorporating global information during inference** and **additional optimization of the variational parameters** tightens the lower bound and helps learning better generative models
- Further work needed to study why additional optimization helps learning models in sparse data compared to dense data (Hjelm et. al)
- Rich variety of information stored in the Jacobian of the conditional distributions:
- The log-singular values directly yield information on the number of dimensions utilized
- The rows of the Jacobian matrix may be used as the learned representation under the model