A Notation

$n \in \{1 \dots N\}$ $m \in \{1 \dots M\}$ $\ell \in \{1 \dots L\}$	number of data points n number of inducing variables m number of output dimension of the GP ℓ
\mathbf{w}_n \mathbf{x}_n \mathbf{y}_n	the latent variable corresponding to data point $n, \mathbf{w} \in \mathbb{R}^{D_{\mathbf{w}}}$ n^{th} input variable, $\mathbf{x} \in \mathbb{R}^{D_{\mathbf{x}}}$ n^{th} output variable, $\mathbf{y} \in \mathbb{R}^{D_{\mathbf{y}}}$
W X Y	matrix collecting all the latent variables $\mathbf{W} = {\{\mathbf{w}_n\}_{n=1}^N, \mathbf{W} \in \mathbb{R}^{N \times D_{\mathbf{w}}}}$ matrix collecting all the inputs $\mathbf{X} = {\{\mathbf{x}_n\}_{n=1}^N, \mathbf{X} \in \mathbb{R}^{N \times D_{\mathbf{x}}}}$ matrix collecting all the observations $\mathbf{Y} = {\{\mathbf{y}_n\}_{n=1}^N, \mathbf{Y} \in \mathbb{R}^{N \times D_{\mathbf{y}}}}$
σ^2	observation noise
$egin{array}{l} {f f}_\ell \ {m f}(\cdot) \ k(\cdot,\cdot) \end{array}$	function $f_{\ell}(\cdot)$ evaluated at certain inputs collection of GPs, $\{f_{\ell}(\cdot)\}_{\ell=1}^{L}$ prior covariance function of the d^{th} GP
$egin{array}{c} \mathbf{Z} & & \ \mathbf{u}_\ell & \ \mathbf{U} & & \ \mathbf{m}_\ell & \ \mathbf{S}_\ell & \end{array}$	locations of variational pseudo-inputs evaluations of the ℓ^{th} GP at the pseudo-inputs: $\mathbf{u}_{\ell} = \{f_{\ell}(\mathbf{z}_m)\}_{m=1}^{M}$. collection: $\mathbf{U} = \{\mathbf{u}_{\ell}\}_{\ell=1}^{L}$ variational posterior mean \mathbf{u}_{ℓ} variational posterior covariance of \mathbf{u}_{ℓ}

Table 3: Nomenclature

B Network Architectures

In all neural network setups we apply dropout to the output of the hidden layers. The optimal dropout rate is found using grid-search over $\{0.2, 0.1, 0.01, 0.0\}$. We use the Adam optimizer for optimization and perform grid-search over $\{0.01, 0.001, 0.001\}$ to determine the optimal learning rate. The biases are initialized to zero and the weights using the Xavier distribution [14].

Heteroscedastic noise modeling on UCI datasets Given the modest size of the UCI datasets we choose a relatively small encoder and decoder architecture. The encoder has layers of the following size: $D_{\mathbf{x}} + D_{\mathbf{y}}$, 50, 100, 50 and $2 \times D_{\mathbf{w}}$. The decoder layers have size $D_{\mathbf{x}} + D_{\mathbf{w}}$, 10, 50, 50, 10 and $D_{\mathbf{y}}$. For this experiment we choose a unidimensional latent variable, $D_{\mathbf{w}} = 1$. The targets of the UCI datasets are also 1D. We use the tanh activation function for all hidden layers and the linear activation function for the final layer.

Density estimation on MNIST In this experiment we follow Wu et al. [29] for the network architecture. The decoder has 5 fully connected layers of size: $D_x + D_w$, 64, 256, 256, 256, 1024 and D_y . The tanh activation function is applied to the outputs of the hidden layers and the sigmoid function on the final one. The encoder's fully connected layers have size: $D_x + D_y$, 256, 64, $2 \times D_w$. We use the tanh activation function for the encoder's hidden layers and the linear activation function for the final layer. The inputs are one-hot encoding of the labels, $D_x = 10$ and $D_y = 784$.

C Inference details

In section 3.2 we follow Hensman et al. [15] and use sparse GPs to approximate the full GP. We do this by introducing M inducing or pseudo-inputs $\mathbf{z}_m \in \mathbb{R}^{D_{\mathbf{x}}+D_{\mathbf{w}}}$ and collect them in the matrix $\mathbf{Z} = \{\mathbf{x}_m\}_{m=1}^M$. The inducing outputs $\mathbf{u}_{\ell} = f_{\ell}(Z)$ are the function $f_{\ell}(\cdot)$ evaluated at the inducing inputs. Similarly, we collect the inducing outputs over all dimensions in the matrix $\mathbf{U} = \{\mathbf{u}_{\ell}\}_{\ell=1}^L$. We assume a Gaussian prior for the inducing outputs $p(\mathbf{u}_{\ell}) = \mathcal{N}(\mathbf{0}, \mathbf{K})$ and define a posterior distribution of the form $q(\mathbf{u}_{\ell}) = \mathcal{N}(\mathbf{m}_{\ell}, \mathbf{S}_{\ell})$. The sparse GP framework assumes \mathbf{u}_{ℓ} and \mathbf{f}_{ℓ} to be jointly Gaussian. We can now write $f_{\ell}(\cdot)$ conditioned on the inducing outputs \mathbf{u}_{ℓ} as

$$f_{\ell}(\cdot) \mid \mathbf{u}_{\ell} \sim \mathcal{GP}\left(\mathbf{k}_{\mathbf{Z}}^{\top}(\cdot)\mathbf{K}_{\mathbf{ZZ}}^{-1}\mathbf{u}_{\ell}, k(\cdot, \cdot) - \mathbf{k}_{\mathbf{Z}}^{\top}(\cdot)\mathbf{K}_{\mathbf{ZZ}}^{-1}\mathbf{k}_{\mathbf{Z}}(\cdot)\right),$$

where $[\mathbf{k}_{Z}(\cdot)]_{m} = k(\cdot, \mathbf{z}_{m})$ and $[\mathbf{K}_{ZZ}]_{mm'} = k(\mathbf{z}_{m}, \mathbf{z}_{m'})$. Marginalizing with respect to $\mathbf{u}_{\ell} \sim \mathcal{N}(\mathbf{m}_{\ell}, \mathbf{S}_{\ell})$ leads to the variational posterior $q(f_{\ell}(\cdot))$

$$f_{\ell}(\cdot) \sim \mathcal{GP}\left(\mathbf{k}_{\mathsf{Z}}^{\top}(\cdot)\mathbf{K}_{\mathsf{ZZ}}^{-1}\mathbf{m}_{\ell}, k(\cdot, \cdot) - \mathbf{k}_{\mathsf{Z}}^{\top}(\cdot)\mathbf{K}_{\mathsf{ZZ}}^{-1}(\mathbf{K}_{\mathsf{ZZ}} - \mathbf{S}_{\ell})\mathbf{K}_{\mathsf{ZZ}}^{-1}\mathbf{k}_{\mathsf{Z}}(\cdot)\right) =: \mathcal{GP}(\mu_{\ell}(\cdot), \sigma_{\ell}^{2}(\cdot)).$$

As both, the variational posterior over the GP and the likelihood, are Gaussian, $\mathcal{L}_{\mathbf{w}_n}$ (defined in (3)) can be calculated analytically. We start by using the fact that the likelihood factorizes over the output dimension $\mathbb{E}_{q(f(\cdot))} \left[\log p(\mathbf{y}_n | f(\cdot), \mathbf{x}_n, \mathbf{w}_n) \right] = \sum_{\ell} \mathbb{E}_{q(f_{\ell}(\cdot))} \left[\log p(\mathbf{y}_{n,\ell} | f_{\ell}(\cdot), \mathbf{x}_n, \mathbf{w}_n) \right]$. We define a single term of the previous sum as $\mathcal{L}_{\mathbf{w}_n}^{\ell}$, which can be computed as

$$\begin{aligned} \mathcal{L}_{\mathbf{w}_{n}}^{\ell} &= \mathbb{E}_{q(f_{\ell}(\cdot))} \left[-\frac{1}{2} \log(2\pi\sigma^{2}) - \frac{1}{2\sigma^{2}} \left(y_{n,\ell}^{2} + \mathbf{f}_{n,\ell}^{2} - 2 \, y_{n,\ell} \mathbf{f}_{n,\ell} \right) \right] \\ &= -\frac{1}{2} \log(2\pi\sigma^{2}) - \frac{1}{2\sigma^{2}} \left(y_{n,\ell}^{2} + (\sigma_{\ell}^{2}([\mathbf{x}_{n}, \mathbf{w}_{n}]) - \mu_{\ell}^{2}([\mathbf{x}_{n}, \mathbf{w}_{n}])) - 2 \, y_{n,\ell} \mu_{\ell}([\mathbf{x}_{n}, \mathbf{w}_{n}]) \right). \end{aligned}$$

D Proof of optimality for latent variable posterior

In Section 3.2 we argue that using the optimal free-form distribution $q(\mathbf{w}_n)$ leads to a tighter lower bound. The bound for the optimal and Gaussian form $q(\mathbf{w}_n)$ are, respectively

$$\mathcal{L}_{\text{GAUSS}} = \sum_{n} \left\{ \mathbb{E}_{q(\mathbf{w}_{n})} \mathcal{L}_{\mathbf{w}_{n}} - \text{KL}\left[p(\mathbf{w}_{n}) \| p(\mathbf{w}_{n})\right] \right\} - \sum_{\ell} \text{KL}\left[q(\mathbf{u}_{\ell}) \| p(\mathbf{u}_{\ell})\right]$$
$$\mathcal{L}_{\text{FREE}} = \log \mathbb{E}_{p(\mathbf{W})} \left\{ \exp\left(\sum_{n} \mathcal{L}_{\mathbf{w}_{n}} - \sum_{\ell} \text{KL}\left[q(\mathbf{u}_{\ell}) \| p(\mathbf{u}_{\ell})\right]\right) \right\}.$$

Starting from \mathcal{L}_{FREE} and using Jensen's inequality, we get

$$\mathcal{L}_{\text{FREE}} = \log \mathbb{E}_{q(\mathbf{W})} \left\{ \frac{p(\mathbf{W})}{q(\mathbf{W})} \exp \left(\sum_{n} \mathcal{L}_{\mathbf{w}_{n}} - \sum_{\ell} \text{KL}\left[q(\mathbf{u}_{\ell}) \| p(\mathbf{u}_{\ell})\right] \right) \right\}$$
$$\geq \sum_{n} \left\{ \mathbb{E}_{q(\mathbf{w}_{n})} \mathcal{L}_{\mathbf{w}_{n}} - \text{KL}\left[p(\mathbf{w}_{n}) \| p(\mathbf{w}_{n})\right] \right\} - \sum_{\ell} \text{KL}\left[q(\mathbf{u}_{\ell}) \| p(\mathbf{u}_{\ell})\right].$$

Therefore

$$\log p(\mathbf{Y} \mid \mathbf{X}) \geq \mathcal{L}_{\text{FREE}} \geq \mathcal{L}_{\text{GAUSS}}.$$

E Density estimation on MNIST: complete table

			CVAE		GP-CDE								
	Fixe	d σ^2	σ^2	² optimize	d	Fixe	d σ^2	σ^2 optimized					
N	Test Train		Test	Train	σ_{opt}^2	Test	Train	Test	Train	σ_{opt}^2			
2	-129.72	180.97	-1296.63	956.39	0.01378	161.9	242.2	74.01	130.4	0.0303			
4	-60.03	178.22	-759.18	956.26	0.01364	195.2	254.2	86.59	160.3	0.0310			
8	-31.37	176.76	-616.474	949.83	0.01358	234.1	269.8	124.6	168.2	0.0312			
16	-15.50	173.34	-497.15	924.11	0.01382	290.1	305.8	141.5	163.5	0.0303			
32	-1.04	161.10	-382.20	800.03	0.01577	452.7	443	131.9	131.7	0.0322			
64	18.03	130.78	-227.06	530.40	0.02348	545.2	515.7	141.2	114	0.0342			
128	41.94	97.23	76.15	399.78	0.02959	508.3	447	93.6	89.1	0.0364			
256	52.17	76.18	218.08	325.72	0.03272	606.2	545	108.1	105.4	0.0378			
512	54.48	65.30	244.88	286.38	0.03407	606.7	512	124.2	120.7	0.0388			

Table 4: Log-likelihoods of the CVAE and GP-CDE model. N is the number of training images per class. Higher test log-likelihood is better.

F Additional Figures



Figure 5: Kernel density estimation models, unconditional (top row), and conditioned on the 50 nearest neighbors.



Figure 6: MDN. Top row: 5 components. Bottom row: 50 components

ല É CL) CL Ł b $\gamma_{\rm H}$ v.O Ŭ. 7 Ł 2 Λ à 血 2 11. 12.13 1.3 12 U1 প T 1 η \mathbf{v} 14 υ H m Δ À d 0 4 \square Þ o 5 0 A. α 1 \$ \mathcal{U} வ 📇 41 $\overline{\lambda}_{D}$ h -123 0 --5 C7 (C) 3 8 C77 1223 भिष्टि 🔛 🗰 œ J В Э đ ഹ 100 5 Ð 4 ÷ mo. ਅ੍ਰੀ 17 H て Б F ഷം 65 4.1 540 4D ىن ج 14 0 Ş • E Ē 0 Ľ. బ 30 1.3 1... 8.1 21 10 6-C E T h ሢ 114 GIT. ন্ত 1.1 K ÷÷ 1 1 151 2 2 2 2 ęЗ --1 τa) Q. d d d -5 6 5 CL. -5 ŧ Ŧ -Ŧ \mathbf{f} φ đ đ đ ŧ Ľ ₽ П ট 5 古 5 5 5 Σ 1 ປ 1 D υ \sum 37 1.6 18 ч W 1.1 3. Ħ P Y U वु Ŧ ЭĻ Ħ 1 1 34 D з. ਸ л v 14 Ta 14 ਰ C. 1 10 R 410 ų Ľ ባብ 😡 ъ ц., (£3), ซ С --_ _ _ **/ _ _** CT. a. 6 b 3 P iΦ. 10 ęχ В П F আ ۲, b 6.91 꿥 Ŧ 2 P ਰਿ -56 5 -5 ズ C) n J ى 3 1 ł Т Θ 1 <u>1 m</u> tan) ব্য t 6.0 1997 ਸ਼ Ð \mathcal{T}_{a} Ū. 1 = . Đ Ð ÷. Æ $\overline{\mathcal{N}}$ 13 E m (1) ch. 114 CE. (D £ E E 문 -F E Н Ŧ .F Ŧ đ C C D ÷ C w 1223 12.3 ÷t. 20 V 10 V (Find 11 -1 100 1 H. -1 Ð Th: -11 8 R 6 E-H h L Ľ \mathcal{H} H ઞ 41 4 ബ് 1 123 Ŧ H H -11 34 de D1 111 100 -12 D ػ 1 đ 1 1 15 ٦ T ÈX. \mathbf{T} J. ₽ Ň Ø P Ð 7 A ρ - 16.0 7 -У P 3 3 1 H 用 12 10 1 -LU. з H 12 v3 3 ₹ ÿ Ð Ð B 35 D 12 Θ Θ T. **q**~ 7 1 14 ۲ $\mathbf{\bar{x}}$ Ŧ

Figure 7: Five samples from the posterior conditioned on 120 labels randomly chosen from the **training** set. The model saw 20 examples of each image. An example image from the data is shown to the left of each set

rt	ЧŤ.	1	Ŧ.		ŧ	74	F	Ħ	F	臣	F	ħ	1		1	*	-5	ਜ	Ħ	184	Ŧ	म	đ
സ	1927	tent	Cal	Se.	예	ದ	с ц	¢	-	dit i	ċΩ	Ц	125	Ē	12	Ð	E.	Ф ₁	ä		h	h	h
Þ	Ь	12	b	b	Ь	ഹ	न्द्री	eg)	(d)	ref.	ej1	∦ a		40		. He	Þ	- ?	2	2	2	P	2
ふ	A	Ē.	A	4	Δ	Ŷ	Ц	66	-51	ē,	1	ম্প	G	М	P	(H)	ল	9	2	Ø	0	Ø	(7)
କ	8	6	6	8	8	ጽ	M	8	2	莆	æ	Þà	邗	1	11	P	博	ъJ	¢0	-	H	¢9.,	e,i
fu	Ħ		H			γp	(P)	æ	UP!	w	8	વ	-	T	E	Ħ	Ħ	ಕಿ	0	4	đ	4	4
XVVS	3	3	3	3	ø	2	2	Ð	$\langle \overline{c} \rangle$	7	÷.	6	6	6	6	6	6	ର	R	P.	6	Ø,	9
තු	0		(15)	60	600	A	15	đ	Ð	ð	3	<u> </u>	E .	0	-	æ,	125	ఓం	te i	62	Ð	ы	Ο.
0°	Û	6 79		80	ŧ.	L	£	r	185) 1851	Ф	4	ਖ	¥	Ħ	PH.	E.	d.	ਫ	ũ.	Œ	đ	1	IJ
Ψ	Ч.	H	1	闸	H	э.	ϕ_{i}	ð.,	¢Å,	Ø.	5 .	పి	U	Ð	Ð	P.	P	4	t)		ŧ.	t.	
M	0	9	-	9	6	₿.	4	-	0	8	õ	ᡉ	b,	ø	(iii)	-	Ø	N	M	Ŧ	Ħ	H.	B
म	ল	म	.	E.	Ŧ	P	e	P	Ē	Ρ	-	H	睫	4	建築	₽	ŧ	Ч	謿	닅	P	Ð,	12
ರ	02.	η	3	1	đ	벁	G 1	曲	W.	Ф,	141	o† I	1	ø	-	EX.	1	าโไ	¢,	ιD)	ф	фĬ,	ø
\mathcal{F}	2	≥ ai	2	1	15	И	Ħ.	1	Ð	1	١ <u>n</u>	0	۲	٥	Φ	ō,	٩	$\left \mathbf{Q} \right $	(E)		ଙ୍କ	2	ଙ୍କ
4	E	E	1			<i>8</i> 6	1	£.	靐	m	(CC	ક્ષ	¥8.	1	#1	W	橋	12	۵	Ð	ŧ.	Ð	0
0	0	G	0	Ø	0	ಜ		an)	Ħ	۵,	25	\mathcal{V}	ы	Li,	(E)	Ð	ω	1	1		Ŧ	-	đ.
Ъ	1	T.	1	电	F	F	4		μų.	-	-	v	M.	Ð	1	म	W	٤3	U.	B	tit.	ы	ы
ЪС,	a r	đ	æ	9	æ	ಹ	E	Ċ.	ट	Ø	đ		63	411	<u>en</u>	itt.	65.	(ä	3		3	麋	3
л	ы	म	(64 2)		Ħ.	ŕĈ	Ð	1		1	(E)	<i>⊭1</i> 1	1	1	tirt.	th.	93	വ	T	ç ų	D.)	¢4,	Ġ.
f	E.	Ŧ	T	1	Æ	٩	٥		4	-	•	$\overline{\mathcal{V}}$	<u>a</u>	۵I)	Ŧ	Ū,		ß	đ	et.i	¢.	#	±1
7	币	5	E		4	$\langle \nabla \rangle$	P	P	2	P	- 19	₹.	T	$\underline{\mathbf{Z}}$	5	1	5	a	0	G		0	13
Υ	14	2	T	Ľ	Ý	ଜ୍		9	R	3		\mathcal{V}_{\parallel}	Ø	U	$\Theta^{\rm e}$	${\boldsymbol{\mu}}$	Ð	0	1	0	eth)	0	P
Ē	Ċ.	E.	es.	C .	¢.	Ð	ы	Ě	E.	1	10	q	뚶	8	Ħ	늰		\asymp	H,	70	E.	1	N.
Ł	1	1	1	2	¥.	8	-	0	199	40	ab	ন্ত্র	8	Ē	R	Ħ	B	뮰	E	E			
ਸਾ	th.	13	(H)	Ш.	Ħ	\neq		P	2	1	1	ପ	Ð	8	0	Ð	8	h	Đ	Ŀ	b	1	÷.
ю	E.	٥	۵	ñ	Þ	ন্দ্র	eta.	Ū.	hill	-	140	ዋ	-	Φ	Φ	ų,	4	T	Ť	Ŧ	T	<u>en</u>	Ŧ
ሌ	(†)	63	÷.	¢.	۰	YE		Ħ	6	Ø	E	Б	6	Ø	£	£	Ø	भ	ų.		M	4	4
\Box	Ø	Ð	100	E.	(E)	ν		14	140	de	u	8	*	1	6	3	Ċ	ഇ	B r	1	ø	B	
ഊ	UI.	(C))	a)	(U)		শ	म	-	T	-	1	4	\$	1	Ф	4	ф	В		-	H.	E.	6
냉	H.		6	W	₩.	Ŝ	E.	2	œ	-	8	え		ø	63	12	10	X	ज	Ψ.	퐦	-	ų

Figure 8: Five samples from the posterior conditioned on 120 labels randomly chosen from the **test** set. The model saw 4 examples of each image. An example image from the data is shown to the left of each set



Figure 9: The mean of the GP mapping conditioned on the 1D latent variable, for a dataset of 100 '1's. Top: the analytic solution of Titsias and Lawrence [26]. Middle: our natural gradient approach, with a step size of 0.1. Bottom: using the Adam optimizer to optimize the variational parameters m and S. We see that the ordinary gradient approach is prone to getting 'stuck' in poor local optima, due to the difficulty of optimization.



Figure 10: The training objective for the dataset of 100 '1's. The three plots show the same three curves with different ranges on the y-axis to highlight the similarities and differences. We can see that natural gradients provide a striking improvement. See Fig. 9 for the samples at the end of training