A Deep Generative Model for Interactive Data Annotation through Direct Manipulation in Latent Space

Hannes Kath[1,2], Thiago S. Gouvêa[1], and Daniel Sonntag[1,2]

[1] German Research Center for Artificial Intelligence (DFKI), Oldenburg, Germany
{hannes_berthold.kath, thiago.gouvea, daniel.sonntag}@dfki.de
[2] University of Oldenburg, Applied Artificial Intelligence (AAI), Oldenburg, Germany

Abstract. The impact of machine learning (ML) in many fields of application is constrained by lack of annotated data. Among existing tools for ML-assisted data annotation, one little explored tool type relies on an analogy between the coordinates of a graphical user interface and the latent space of a neural network for interaction through direct manipulation. In the present work, we 1) expand the paradigm by proposing two new analogies: time and force as reflecting iterations and gradients of network training; 2) propose a network model for learning a compact graphical representation of the data that takes into account both its internal structure and user provided annotations; and 3) investigate the impact of model hyperparameters on the learned graphical representations of the data, identifying candidate model variants for a future user study.

Keywords: Deep Generative Model · Self-supervised Learning · Variational Autoencoder

1 Introduction

Modern machine learning (ML) algorithms such as deep neural networks (DNNs) have shown outstanding performance on various tasks, especially on supervised learning. Still, the impact of ML in many application areas is limited by the low availability of annotated data, as well as by technical entry barriers for domain experts, leading to a cross-domain need for user-friendly tools to annotate vast datasets. To address this need, we propose and investigate a deep generative model (DGM) for the development of a data annotation tool that visualises data samples in a 2D or 3D graphical user interface (GUI).

Previous work has shown that interactive GUIs using the latent space of ML systems can facilitate data interaction [14]. The authors show that ML systems can learn representations of a data domain from which users can obtain actionable views using GUIs, such as virtual reality (VR). Another tool for ML-assisted data annotation operating on latent representations is Seadash [7]. Seadash is a tool for graphical data programming [15] that operates based on redundant, parallel data transformations.

In this paper, we take a general approach to improving data efficiency, which is to first learn a representation that captures the data structure [4], then annotate some data points, and finally restructure the data incorporating new annotated data points, facilitating further annotation. A powerful method for learning useful representations in an annotation-free manner is self-supervised learning (SSL) [11, 12]. In SSL, a model learns to predict labels that are automatically computed from the input data itself, e.g. by reconstructing input. The learned representations are then fine tuned based on the class annotations provided by the user. To exploit the limited amount of annotated data, we use a DGM, following the promising approach of Kingma et al. [9]. SSL can be considered as a special task in generative models that involves imputing missing data as a means to improve classification accuracy [9]. DGMs learn entangled representations, which means that the dimensions of the latent space do not always have a consistent meaning along the axes [13]. Using a DGM that learns disentanglement [13] in such a way that the decision boundaries or the respective clusters are clearly visible to the user, we structure the latent space.

2   Application

The DGM described in this paper is designed to run a data annotation system with 2D or 3D representations. The overall system architecture is shown in figure 1a. Leading to the model requirements outlined in section 3, this section provides a concise overview of the user's workflow. Details about the development and presentation of the application can be found in separate publications for a 2D GUI [2] and a 3D GUI [3].

The workflow of the application follows the concepts of direct manipulation (DM), which focus on physical, incremental and reversible user interactions on displayed objects of interest whose effects are immediately visible [16]: A partially labelled dataset is embedded in a latent 2D or 3D space represented in the GUI. Similar inputs are clustered so that multiple data samples can be annotated at once (e.g. with the Lasso tool). New annotated data samples are used to re-train the DGM, which triggers movement of the data samples: data samples of the same class are pulled together while different classes are pushed apart, creating the illusion of an invisible force like magnetism.

By applying the concept of DM, we derive three metaphors: space (latent space represented as interface space), time (intermediate training results represented as movement) and force (backpropagation represented as physical force). In this work, we therefore focus on the learned embedding of data points, the training process and backpropagation of the DGM, while the user interactions in terms of space, time and force take place in the GUI.

3   Model description

Functional requirements for the model derived from the system architecture in figure 1a include learning relevant data structures without annotations, mapping

them into a 2D or 3D latent space, a way for the user to explore, inspect and annotate data, as well as a way to helpfully customize the latent space for the user based on the added annotations. The central non-functional requirement derived from the metaphors described in section 2 is that the user interface needs to reflect essential concepts of the ML system, namely the latent space, the training progress and the backpropagation. Considering these requirements, we developed a DGM architecture (figure 1b) based on a variational autoencoder (VAE) [10] and extended it with a classification header.
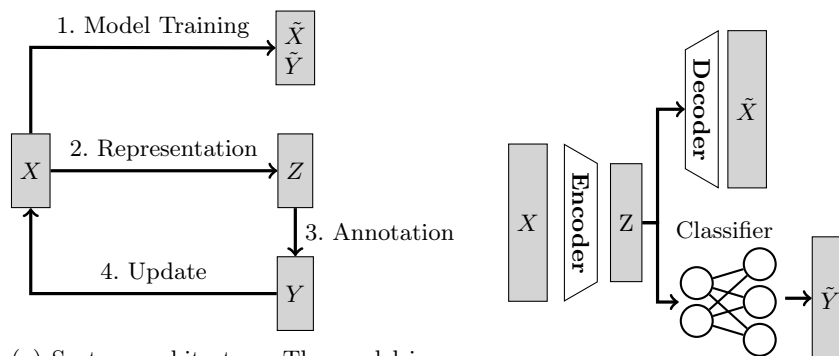
(a) System architecture: The model is trained (1) using the (partially annotated) dataset. Based on the trained model, the dataset is presented in latent space to the user (2), who can annotate samples (3). New annotated data trigger a new system cycle that also uses the new annotations (4).

(b) Model architecture: The raw dataset is processed and represented in latent space (encoder). For training purposes the representation is reconstructed (decoder) and annotated data is presented to the classification header (classifier).

Fig. 1: Presentation of system and model architecture with X (dataset), Z (representation in latent space), $\tilde{X}$ (reconstruction), Y (label annotated by user) and $\tilde{Y}$ (label predicted by classifier).

VAE. Figure 1b shows the VAEs architecture besides the classification header. The encoder $q_\phi(Z \mid X)$ captures data structures from the input data X and represents them in the latent space using the representation Z. The dimensionality of the latent space is chosen low to force the learned representation to capture the most salient features of the training data and to allow the representation of the entire state space in a 2D or 3D environment. The decoder $p_\theta(X \mid Z)$ is used for training purposes and calculates the reconstruction $\tilde{X}$ using Z. The parameters ϕ and θ are learned jointly using gradient descent on random minibatches. A VAE is a probabilistic model, because each input X induces a probability distribution over the entire latent space of variable Z. In the standard implementation of a VAE, Z is parameterized as a multivariate Gaussian whose parameters μ_Z

and Σ_Z are computed deterministically from X, and downstream operations are performed on samples from this distribution [10]. While samples of Z are used in learning, we propose to use the mean μ_Z for visualizations.

Classification head. The latent space of the VAE is displayed to the user for annotation. A user might find it helpful that the space is transformed to group data points of the same category. Since VAEs cannot handle annotations, we introduce the variable Y and a classification head that computes the class labels $\tilde{Y}$ from the latent variable Z. This classification head is a simple multilayer perceptron (MLP).

Loss functions. The VAE and MLP are jointly optimized by minimizing the sum of all loss components:

$$\mathcal{L} = \mathcal{L}_{\mathrm{reconst}}(X, \tilde{X}) + \beta_{KL} D_{KL}\left(q_\phi(Z \mid X) \,\|\, p(Z)\right) + \beta_{\mathrm{classifier}} \mathcal{H}(Y, \tilde{Y}), \quad (1)$$

where the first two terms are from [5], $\mathcal{H}$ is the cross entropy loss between Y and $\tilde{Y}$, and β_{KL} and $\beta_{\mathrm{classifier}}$ serve as the weighting coefficients for the loss terms.

4 Experiments

To analyze the potential of the proposed method for designing the user interface by varying hyperparameters of the ML system, we analyzed their impact on the user experience. As an apparatus, we created a convolutional VAE similar to the one in [5] with a 2D latent space and trained it on $10\,\%$ of the MNIST dataset [6]. Data preprocessing included one hot encoding the labels, which allowed using a cross entropy loss in the classifier. We used the common Adam optimizer with a learning rate of $5 \cdot 10^{-3}$. Our MLP classifier head has 10 units per layer and was trained with $100\,\%$ classified data. We conducted experiments by varying the following hyperparameters:

Hidden layers of classification head. We hypothesized that classification heads with lower capacity would constrain the learned representation to simpler structures and thus force easier separation for the user. To test that, we pretrained the VAE using unsupervised learning (figure 2a, left) and added a classifier head with either two (figure 2a, right) or none (reducing to a logistic regression, figure 2a, center) hidden layers. The results show that self-supervised training already leads to visible clustering. Adding a classification head with logistic regression leads to sharper separation of classes, whereas this is not clearly observed with a higher capacity classification head.

Number of training epochs. For investigating the model stability, after self-supervised pretraining (i.e., the VAE component), we added the logistic regression classification head and trained it for different numbers of epochs (figure 2b). The results show a stable change that can be continuously displayed in a tool to show changes to the user immediately.
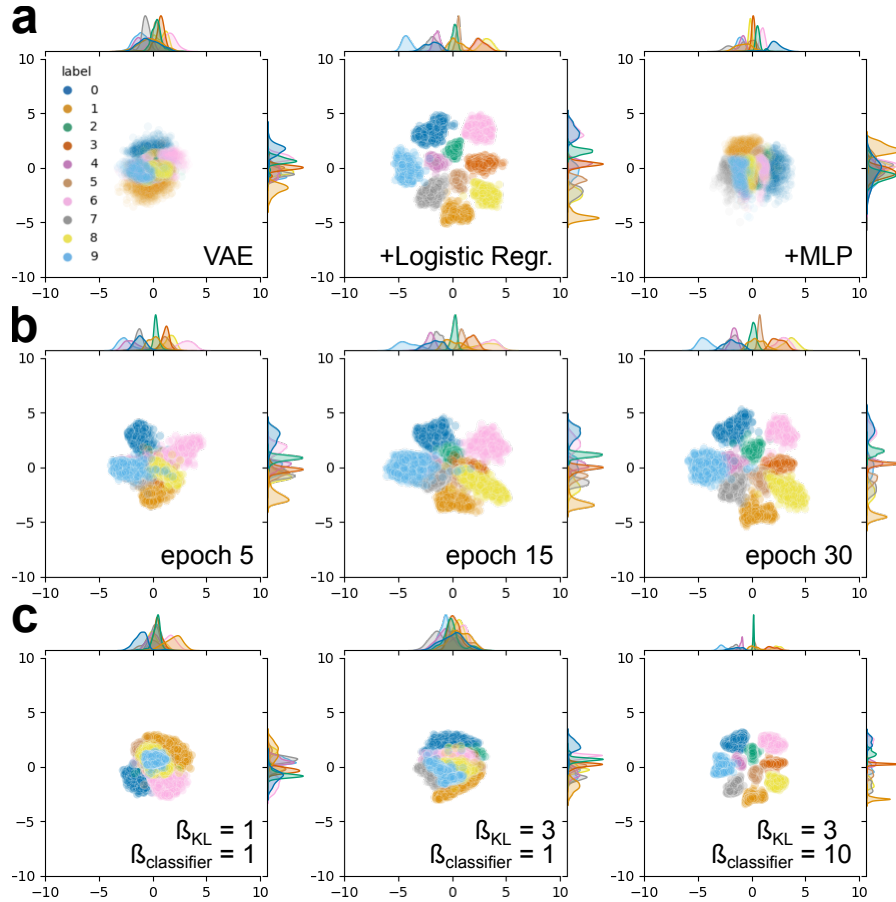
Fig. 2: **Visualisation of MNIST dataset embedded in 2D latent space learned by model variants.** Scatter plots show mean of latent variable Z for each input X; marginal plots are marginal histograms; colours denote label class. **(a)** Different classifier heads. *left:* No classifier; initial condition for all other learned representations. *center:* Logistic regression. *right:* MLP with 2 hidden layers. Hyperparameters set to $\beta_{KL} = 3$, $\beta_{\text{classifier}} = 100$, $n_{\text{epochs}} = 50$. Logistic regression used for all other experiments. **(b)** Learning observed after 5, 15, and 30 training epochs; compare with initial (a, left) and final (a, center) states. **(c)** Different weights applied to loss components. Compare with (a, center).

Components of the loss function. We varied β_{KL} and $\beta_{\text{classifier}}$ in equation (1). The representation learned with neutral values is shown in the figure 2c, left. Figure 2c, right and figure 2a, center show that increasing $\beta_{\text{classifier}}$ leads to learning representations with tighter clusters. Increasing β_{KL} reduces the overall spread of data points (figure 2c, center). This is explained by the choice of prior $p(Z)$ to be a standard Gaussian for each of the dimensions of Z, and the fact that β_{KL} scales the divergence between input representation in latent space and the chosen prior. By choosing even higher values for β_{KL}, the second term of equation (1) gains priority, which leads to a collapse where the encoder outputs $\mu_Z = 0$ and $\Sigma_Z = 1$ for all inputs. The likelihood of this collapse occurring can be controlled by weighing the KL component [5].

5   Conclusions and future work

A real need has been identified for interactive machine learning, shifting the focus from developing more accurate algorithms to improving applicability, which may include other metrics such as productivity and interpretability [1, 17]. We have described a neural architecture that can drive a GUI for interactive representation learning for ML-assisted data annotation, and we have found that the hyperparameters (classification header, number of training epochs, and loss function weights) have potential effects on user experience. Our network enables DM of data in latent space by establishing three analogies relevant to an immersive experience: (latent as interface) space, (training as frame) time, and (stochastic gradient descent as physics engine) force. We are currently exploring different learning rates to produce different degrees of stability of representations, as well as continuous learning methods [8] to handle the incremental availability of labels provided by the interface front-end.

References

1. Amershi, S., Cakmak, M., Knox, W.B., Kulesza, T.: Power to the people: The role of humans in interactive machine learning. Ai Magazine **35**(4), 105–120 (2014)
2. Author: A human-in-the-loop tool for annotating passive acoustic monitoring datasets. In: Proceedings of the 32nd International Joint Conference on Artificial Intelligence, IJCAI (2023)
3. Author: A virtual reality tool for representing, visualizing and updating deep learning models. In: KI 2023 (submitted)
4. Bengio, Y., Courville, A., Vincent, P.: Representation learning: A review and new perspectives. IEEE transactions on pattern analysis and machine intelligence **35**(8), 1798–1828 (2013)
5. Burgess, C.P., Higgins, I., Pal, A., Matthey, L., Watters, N., Desjardins, G., Lerchner, A.: Understanding disentangling in β-vae. arXiv preprint arXiv:1804.03599 (2018)
6. Deng, L.: The MNIST database of handwritten digit images for machine learning research [best of the web]. IEEE Signal Process. Mag. **29**(6), 141–142 (2012). https://doi.org/10.1109/MSP.2012.2211477, https://doi.org/10.1109/MSP.2012.2211477

7. Gouvêa, T.S., Troshani, I., Herrlich, M., Sonntag, D.: Annotating Sound Events Through Interactive Design of Interpretable Features. HHAI2022: Augmenting Human Intellect pp. 305–306 (2022). https://doi.org/10.3233/FAIA220225, https://ebooks.iospress.nl/doi/10.3233/FAIA220225

8. Hadsell, R., Rao, D., Rusu, A.A., Pascanu, R.: Embracing change: Continual learning in deep neural networks. Trends in cognitive sciences **24**(12), 1028–1040 (2020)

9. Kingma, D.P., Mohamed, S., Rezende, D.J., Welling, M.: Semi-supervised learning with deep generative models. In: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D., Weinberger, K.Q. (eds.) Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada. pp. 3581–3589 (2014), https://proceedings.neurips.cc/paper/2014/hash/d523773c6b194f37b938d340d5d02232-Abstract.html

10. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013)

11. LeCun, Y., Misra, I.: Self-supervised learning: The dark matter of intelligence (Mar 2021), https://ai.facebook.com/blog/self-supervised-learning-the-dark-matter-of-intelligence/

12. Liu, X., Zhang, F., Hou, Z., Mian, L., Wang, Z., Zhang, J., Tang, J.: Self-supervised learning: Generative or contrastive. IEEE Transactions on Knowledge and Data Engineering (2021)

13. Paige, B., Siddharth, N., van de Meent, J., Desmaison, A., Goodman, N.D., Kohli, P., Wood, F.D., Torr, P.H.S.: Learning disentangled representations with semi-supervised deep generative models. In: Guyon, I., von Luxburg, U., Bengio, S., Wallach, H.M., Fergus, R., Vishwanathan, S.V.N., Garnett, R. (eds.) Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA. pp. 5925–5935 (2017), https://proceedings.neurips.cc/paper/2017/hash/9cb9ed4f35cf7c2f295cc2bc6f732a84-Abstract.html

14. Prange, A., Sonntag, D.: A demonstrator for interactive image clustering and fine-tuning neural networks in virtual reality. In: German Conference on Artificial Intelligence (Künstliche Intelligenz). pp. 194–203. Springer (2021)

15. Ratner, A.J., De Sa, C.M., Wu, S., Selsam, D., Ré, C.: Data programming: Creating large training sets, quickly. Advances in neural information processing systems **29** (2016)

16. Shneiderman, B.: Direct manipulation for comprehensible, predictable and controllable user interfaces. In: Moore, J.D., Edmonds, E.A., Puerta, A.R. (eds.) Proceedings of the 2nd International Conference on Intelligent User Interfaces, IUI 1997, Orlando, Florida, USA, January 6-9, 1997. pp. 33–39. ACM (1997). https://doi.org/10.1145/238218.238281, https://doi.org/10.1145/238218.238281

17. Simard, P.Y., Amershi, S., Chickering, D.M., Pelton, A.E., Ghorashi, S., Meek, C., Ramos, G., Suh, J., Verwey, J., Wang, M., et al.: Machine teaching: A new paradigm for building machine learning systems. arXiv preprint arXiv:1707.06742 (2017)