

EXPLAINING LATENT REPRESENTATIONS OF GENERATIVE MODELS WITH LARGE MULTIMODAL MODELS

Mengdan Zhu¹, Zhenke Liu¹, Bo Pan¹, Abhinav Angirekula², Liang Zhao¹

¹Department of Computer Science, Emory University

²Department of Computer Science, University of Illinois Urbana-Champaign
 {mengdan.zhu, zhenke.liu, bo.pan, liang.zhao}@emory.edu,
 aa125@illinois.edu

ABSTRACT

Learning interpretable representations of data generative latent factors is an important topic for the development of artificial intelligence. With the rise of the large multimodal model, it can align images with text to generate answers. In this work, we propose a framework to comprehensively explain each latent variable in the generative models using a large multimodal model. We further measure the uncertainty of our generated explanations, quantitatively evaluate the performance of explanation generation among multiple large multimodal models, and qualitatively visualize the variations of each latent variable to learn the disentanglement effects of different generative models on explanations. Finally, we discuss the explanatory capabilities and limitations of state-of-the-art large multimodal models.

1 INTRODUCTION

Latent variable based data generation has emerged as a state-of-the-art (SOTA) approach in the field of generative modeling (Mittal et al., 2023; Deja et al., 2023; Patil et al., 2022). This technique leverages latent variables to learn underlying data distributions effectively and also generate high-quality samples (Vahdat et al., 2021). One of the key advantages of using latent variables is their ability to capture the underlying structure in high-dimensional data. However, understanding and interpreting such latent variables is challenging and often requires human expertise for meaningful interpretation. Learning interpretable representations of the data generative latent factors is an important topic for the development of artificial intelligence that is able to learn and reason the same as humans do (Higgins et al., 2016). Large multimodal models (LMMs) have accomplished remarkable progress in recent years (Wang et al., 2024; Wu et al., 2023; Yin et al., 2023; Bai et al., 2024; Ling et al., 2023). LMM is more similar to the way humans perceive the world (Yin et al., 2023). We thus consider using large multimodal models to automatically explain the latent representations.

Recently, one of the most powerful LMMs is the instruction-following LMM (Li, 2023). LLaVA and InstructBLIP are two instruction-following LMMs that achieve SOTA performance on many datasets. Instruction-following models use instruction tuning to enhance their abilities to understand and follow human-given instructions. Instruction tuning involves further fine tuning LLMs using $\langle instruction, response \rangle$ pairs to better align human intent with model behavior (Wang et al., 2024). LLaVA was introduced in the paper Visual Instruction Tuning (Liu et al., 2024), and then further improved in Improved Baselines with Visual Instruction Tuning (referred to as LLaVA-1.5) (Liu et al., 2023). Likewise, InstructBLIP (Dai et al., 2024) is a large multimodal model that adds instruction tuning on the basis of its previous version BLIP-2 (Li et al., 2023). Google Bard is a conversational AI service developed by Google, initially powered by LaMDA with a range of models to follow. In this work, we propose a framework to comprehensively explain each latent variable in generative models and evaluate the performance of explanation generation of GPT-4-vision with several popular LMMs: Google Bard, LLaVA-1.5, and InstructBLIP. To the best of our knowledge, we are the first to use LMM to explain the latent representations of the generative models.

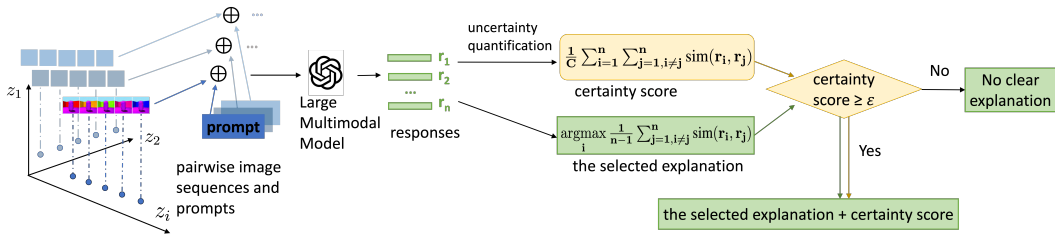


Figure 1: The model framework consists of generating an image sequence with a progressively varying latent variable, combining it with a prompt to pass to a large multimodal model to provide some response samples and finally utilizing an uncertainty measure to select an explanation for that specific latent variable and decide whether there is a clear explanation to display.

2 METHODS

Problem Formulation We start from a dataset $\mathcal{D} = \{X, Z\}$ where the images $\mathbf{x} \in \mathbb{R}^N$ and the data generative latent variables $\mathbf{z} \in \mathbb{R}^M$. We then can train a generative model that learns the joint distribution of the data \mathbf{x} and latent variables \mathbf{z} . Here, our goal is to explain each latent variable z_i individually, where i ranges from 1 to the number of latent dimensions, and quantify the uncertainty of the explanation to make sure the explanations presented are reliable and responsible.

In this work, we confront two primary challenges associated with the representation and interpretation of the latent space. Firstly, the latent space is difficult to explicitly represent. To address this, we interpolate a specific latent dimension z_i at a time and subsequently decode these into an image sequence. This method allows us to visualize the latent representations. Secondly, not all latent variables have semantic meaning. To tackle this issue, we introduce an uncertainty approach to distinguish between interpretable and uninterpretable latent variables.

Framework To start with, we trained three generative models for each dataset. After the generative models are trained, a latent representation \mathbf{z} is first sampled from an isotropic unit Gaussian distribution $\mathcal{N}(\mathbf{0}, I)$. Given a certain latent variable z_i , we perturb z_i to observe the possible value range of z_i and keep the other latent variables constant. A sequence of images can then be generated by decoding a series of manipulated latent vectors. As a result, this generated sequence of images shows how a latent variable z_i changes gradually. Further, we pass this image sequence with a prompt to a large multimodal model to explain the latent variable, as well as its changing pattern. Finally, we quantify the uncertainty to decide which response to select as an explanation and whether to display the selected explanation as shown in Figure 1.

Not every latent variable is interpretable, so we need to find an approach to determine whether a latent variable is interpretable. Similarity and entropy are two major ways to measure uncertainty in natural language generation(NLG). Since the likelihood of tokens is not available in GPT-4-vision, we used the measure of similarity in this paper. However, the use of similarity is not a limitation of the uncertainty measure. Other methods, such as predictive entropy, semantic entropy (Kuhn et al., 2023), and $P(\text{true})$ (Kadavath et al., 2022), could also be used to quantify uncertainty if applicable. To measure the uncertainty of the responses of our large multimodal model, GPT-4-vision, we sampled n times from the GPT-4-vision to generate the responses $R = \{r_1, r_2, r_3, \dots, r_n\}$. The certainty score of the explanation is the average similarity of the responses $R : \frac{1}{C} \sum_{i=1}^n \sum_{j=1, i \neq j}^n \text{sim}(r_i, r_j)$, where $C = n * (n - 1)/2$. To find the threshold of the interpretability of the latent variables, we denote the true label of the interpretability of each latent variable as 1 if a human can see a clear pattern in the generated images, otherwise we denote it as 0. The certainty score of the explanation for each row of generated images is its predicted probability. We can then compute their AUC (area under the ROC curve) based on different thresholds and choose the one with the highest AUC as our threshold ϵ . Our final output explanation is the explanation that has the highest mean pairwise similarity with other responses if the certainty score is equal or greater than the threshold ϵ . Otherwise, we will output there is no clear explanation.

3 EXPERIMENTS

Datasets We perform evaluations with three datasets, the MNIST dataset of handwritten digits (Lecun & Cortes, 2010), dSprites dataset of 2D shapes (Matthey et al., 2017) and 3dshapes dataset of 3D shapes (Burgess & Kim, 2018). The dSprites dataset consists of 6 ground truth latent factors. These factors are color(white), shape(square, ellipse, heart), scale(6 values), rotation(40 values), position X(32 values) and position Y(32 values) of a sprite. Similarly, the 3dshapes dataset is generated from 6 ground truth latent factors of floor color(10 values), wall color(10 values), object color(10 values), scale(8 values), shape(4 values), and orientation(15 values). The MNIST dataset consists of grayscale handwritten digits(0 through 9).

Visual pattern generation We train three representative variational autoencoder(VAE) models, the standard VAE (Kingma & Welling, 2014; Rezende et al., 2014), β -VAE (Higgins et al., 2016), and β -TCVAE (Chen et al., 2018), with three aforementioned datasets. For each trained model, we manipulate one latent variable between [-3, 3] at a time while keeping others constant, and the trained decoder of the model can produce a series of images that reflect variations along that specific latent dimension. This process is repeated for each latent variable, resulting in $m \times k$ images, where m is the number of latent variables and k is the number of assigned values for each latent variable. We set 6 latent variables and assign the values as `torch.range(-3, 3, 0.6)`, so $m = 6$ and $k = 11$ here.

Explanation generation We use GPT-4-vision as our explanation generator and compare it with a wide range of other large multimodal models: Google Bard, LLaVA-v1.5-13b (Liu et al., 2023), and InstructBLIP (Dai et al., 2024). For GPT-4-vision, LLaVA-v1.5-13b, and InstructBLIP, we all set `temperature = 1, top-p = 1`. We then pass the generated images of each latent variable along with a prompt to the LLMs to produce an explanation for each latent variable. For human annotations, two annotators provide two explanations with various sentence expressions for each image sequence, so we have four human explanations for each image sequence as references. Furthermore, We evaluate the explanations across these LMMs with human annotations using BLEU(Papineni et al., 2002), ROUGE-L(Lin, 2004), METEOR(Banerjee & Lavie, 2005), and BERTScore F1(Zhang et al., 2019).

Additionally, to compute the certainty scores, we try both the cosine similarity and lexical similarity (Kuhn et al., 2023), where *sim* is cosine similarity and rouge-L respectively in $\frac{1}{C} \sum_{i=1}^n \sum_{j=1, i \neq j}^n \text{sim}(r_i, r_j)$. As Table 1 shows, the measure of cosine similarity has a much higher AUC, which means it can better distinguish if there is a clear pattern. So we choose to use the cosine similarity to generate the certainty score for the explanation. The best AUC in the experiment is 0.9694 and its corresponding threshold ε is 0.7434.

Table 1: The evaluation metrics of the classification of the presence of a discernible trend.

Uncertainty Estimate	AUC	F1-score	Precision	Recall
lexical similarity	0.6898	0.9600	0.9412	0.9796
cosine similarity	0.9694	0.9684	1.0	0.9388

4 RESULTS AND DISCUSSION

Quantitative evaluation for explanation generation We give the same prompt and image sequence to GPT-4-vision and other large multimodal models to generate explanations. Overall, GPT-4-vision outperforms other large multimodal models in their explanatory capability to explain the visual patterns of latent variables in Table 2. As shown in Appendix B, GPT-4-vision is the only LMM that can accurately recognize the handwritten digits in MNIST, explain how the digits change, and have the best overall explanations. Bard can answer the latent variable based on the prompt, but the latent variables and the patterns found are not always accurate. LLaVA and InstructBLIP are not able to respond based on the prompt. More specifically, LLaVA can only describe the images, but it cannot explicitly answer what the latent variable is. InstructBLIP often repeats the task content in the prompt, yet fails to provide the required explanations.

Uncertainty analysis We observe that when there is an evident pattern in the images, the certainty score of the corresponding explanation is likely to be high in Appendix A Figure 3, showing the sampled explanations are more consistent. Conversely, when human can not find a clear pattern

Table 2: Evaluation of large multimodal models on the generated explanations of the latent variables of different generative models

Dataset	VAE Model	LMM	BLEU	ROUGE-L	METEOR	BertScore
3dshapes	VAE	<i>GPT-4-vision</i>	0.051	0.196	0.370	0.875
		<i>Bard</i>	0.047	0.167	0.240	0.842
		<i>LLaVA-1.5</i>	0.000	0.169	0.224	0.864
		<i>InstructBLIP</i>	0.000	0.167	0.212	0.843
	β -VAE	<i>GPT-4-vision</i>	0.056	0.195	0.302	0.868
		<i>Bard</i>	0.000	0.161	0.224	0.850
		<i>LLaVA-1.5</i>	0.024	0.162	0.206	0.857
		<i>InstructBLIP</i>	0.027	0.160	0.205	0.842
	β -TCVAE	<i>GPT-4-vision</i>	0.058	0.203	0.293	0.865
		<i>Bard</i>	0.000	0.189	0.211	0.864
		<i>LLaVA-1.5</i>	0.030	0.181	0.206	0.856
		<i>InstructBLIP</i>	0.062	0.133	0.107	0.846
dsprites	VAE	<i>GPT-4-vision</i>	0.051	0.190	0.291	0.858
		<i>Bard</i>	0.052	0.225	0.197	0.864
		<i>LLaVA-1.5</i>	0.062	0.176	0.219	0.855
		<i>InstructBLIP</i>	0.065	0.195	0.205	0.847
	β -VAE	<i>GPT-4-vision</i>	0.061	0.222	0.282	0.867
		<i>Bard</i>	0.049	0.210	0.233	0.856
		<i>LLaVA-1.5</i>	0.048	0.197	0.194	0.856
		<i>InstructBLIP</i>	0.042	0.190	0.225	0.842
	β -TCVAE	<i>GPT-4-vision</i>	0.051	0.193	0.276	0.858
		<i>Bard</i>	0.028	0.209	0.262	0.848
		<i>LLaVA-1.5</i>	0.052	0.184	0.187	0.857
		<i>InstructBLIP</i>	0.041	0.194	0.249	0.849
MNIST	VAE	<i>GPT-4-vision</i>	0.038	0.181	0.291	0.862
		<i>Bard</i>	0.000	0.203	0.237	0.865
		<i>LLaVA-1.5</i>	0.027	0.162	0.265	0.857
		<i>InstructBLIP</i>	0.000	0.144	0.217	0.843
	β -VAE	<i>GPT-4-vision</i>	0.000	0.201	0.279	0.863
		<i>Bard</i>	0.026	0.184	0.208	0.857
		<i>LLaVA-1.5</i>	0.000	0.201	0.215	0.846
		<i>InstructBLIP</i>	0.000	0.156	0.209	0.841
	β -TCVAE	<i>GPT-4-vision</i>	0.040	0.207	0.288	0.863
		<i>Bard</i>	0.041	0.227	0.240	0.857
		<i>LLaVA-1.5</i>	0.099	0.245	0.287	0.864
		<i>InstructBLIP</i>	0.000	0.197	0.237	0.836

in the images, the certainty score is relatively low in Appendix A Figure 4, indicating the sampled explanations are more diverse.

We quantitatively evaluate if the certainty threshold we find can correctly distinguish whether there is a clear trend in the generated image sequences in Table 1. Table 1 demonstrates an overall excellent performance of the certainty threshold of cosine similarity to identify the cutoff for the presence of a discernible trend. In comparison, the uncertainty estimate of cosine similarity can better distinguish the scenarios when there is no clear trend than the uncertainty estimate of lexical similarity.

Qualitative evaluation for disentanglement When latent factors are entangled in the image sequences, like the ones of VAE and β -VAE in Appendix C, the LMM typically can only identify one of the entangled latent factors in the explanations. Moreover, the greater entanglement of latent factors makes it more challenging to identify changes in all latent factors within an explanation and contributes to an increase in the uncertainty of the explanation, such as the latent factor of β -TCVAE is more disentangled and thus has a higher certainty score than the ones of VAE and β -VAE.

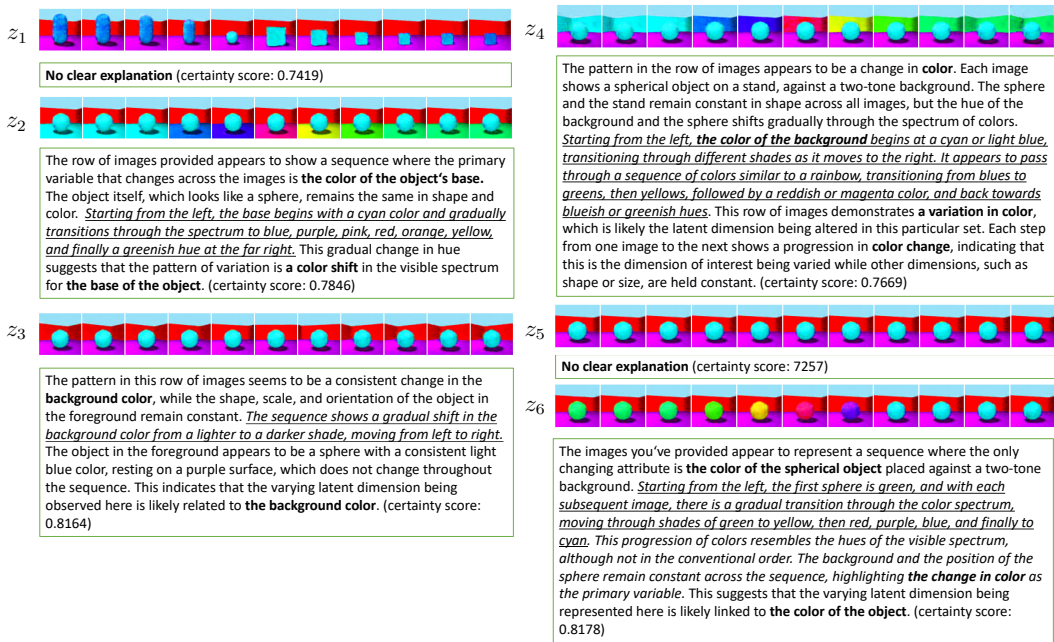


Figure 2: The sample explanations generated by our framework. The latent variables are highlighted in **bold**, and the patterns of the latent variables are in *italics and underlined*.

5 CASE STUDY

Figure 2 shows the explanations generated by our framework for the 3dshapes dataset. If the certainty score of the generated explanation is equal or greater than the certainty threshold ($\epsilon = 0.7434$), it will display the explanation selected. Otherwise, it will return “No clear explanation” like z_1 and z_5 . Our framework can generate an explanation for each latent variable if the latent variable has an evident pattern. The explanation not only can tell what the latent variable is, but also can illustrate how the latent variable changes. Sometimes, the identified latent variable can be wrong, such as it misclassifies the wall orientation as background color in z_3 . This is due to the deficiency in the LMM’s visual capabilities and will be elaborated in the later limitation section.

6 LIMITATION

Although GPT-4-vision is a state-of-the-art large multimodal model and performs the best in the experiments, we still find some of its limitations when evaluating its generated explanations. It is more likely to misinterpret the latent variable as color and is not sensitive to the scale, position, and orientation. Sometimes the description of color patterns is not entirely accurate. Therefore, more work can be undertaken to improve the visual understanding of GPT-4-vision.

7 CONCLUSION

In this work, we propose a framework to comprehensively explain each latent variable in the generative models and visualize the variations of each latent variable. We first analyze the certainty scores of the explanations. The certainty score can tell us if there is a clear trend in the latent variable. We also evaluate and compare the generated explanations with various LMMs. The result shows that GPT-4-vision outperforms other large multimodal models, and we further discuss its limitations and the explanatory capabilities of various LMMs. At last, we evaluate the effects of latent variable disentanglement on the generated explanations. We believe our approach provides an efficient, explainable, and reliable way to learn the latent representations of generative models.

ACKNOWLEDGMENTS

This work is supported by the National Science Foundation (NSF) Grant No. 1755850, No. 1841520, No. 2007716, No. 2007976, No. 1942594, No. 1907805, a Jeffress Memorial Trust Award, Amazon Research Award, Oracle for Research Grant Award, Cisco Faculty Research Award, NVIDIA GPU Grant, Design Knowledge Company (subcontract number: 10827.002.120.04), CIFellowship (2021CIF-Emory-05), and the Department of Homeland Security under Grant No. 17STCIN00001.

REFERENCES

- Guangji Bai, Zheng Chai, Chen Ling, Shiyu Wang, Jiaying Lu, Nan Zhang, Tingwei Shi, Ziyang Yu, Mengdan Zhu, Yifei Zhang, et al. Beyond efficiency: A systematic survey of resource-efficient large language models. *arXiv preprint arXiv:2401.00625*, 2024.
- Satanjeev Banerjee and Alon Lavie. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In Jade Goldstein, Alon Lavie, Chin-Yew Lin, and Clare Voss (eds.), *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pp. 65–72, Ann Arbor, Michigan, June 2005. Association for Computational Linguistics. URL <https://aclanthology.org/W05-0909>.
- Chris Burgess and Hyunjik Kim. 3d shapes dataset. <https://github.com/deepmind/3dshapes-dataset/>, 2018.
- Ricky TQ Chen, Xuechen Li, Roger B Grosse, and David K Duvenaud. Isolating sources of disentanglement in variational autoencoders. *Advances in neural information processing systems*, 31, 2018.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale N Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning. *Advances in Neural Information Processing Systems*, 36, 2024.
- Kamil Deja, Tomasz Trzciński, and Jakub M Tomczak. Learning data representations with joint diffusion models. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 543–559. Springer, 2023.
- Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *International conference on learning representations*, 2016.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, et al. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221*, 2022.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *International conference on learning representations*, 2014.
- Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation. *arXiv preprint arXiv:2302.09664*, 2023.
- Yann LeCun and Corinna Cortes. MNIST handwritten digit database. 2010. URL <http://yann.lecun.com/exdb/mnist/>.
- Chunyu Li. Large multimodal models: Notes on cvpr 2023 tutorial. *arXiv preprint arXiv:2306.14895*, 2023.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pp. 19730–19742. PMLR, 2023.

- Chin-Yew Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pp. 74–81, Barcelona, Spain, July 2004. Association for Computational Linguistics. URL <https://aclanthology.org/W04-1013>.
- Chen Ling, Xujiang Zhao, Jiaying Lu, Chengyuan Deng, Can Zheng, Junxiang Wang, Tanmoy Chowdhury, Yun Li, Hejie Cui, Xuchao Zhang, Tianjiao Zhao, Amit Panalkar, Wei Cheng, Haoyu Wang, Yanchi Liu, Zhengzhang Chen, Haifeng Chen, Chris White, Quanquan Gu, Jian Pei, and Liang Zhao. Domain specialization as the key to make large language models disruptive: A comprehensive survey. *arXiv preprint arXiv:2305.18703*, 2023.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. *arXiv preprint arXiv:2310.03744*, 2023.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36, 2024.
- Loic Matthey, Irina Higgins, Demis Hassabis, and Alexander Lerchner. dsprites: Disentanglement testing sprites dataset. <https://github.com/deepmind/dsprites-dataset/>, 2017.
- Sarthak Mittal, Korbinian Abstreiter, Stefan Bauer, Bernhard Schölkopf, and Arash Mehrjou. Diffusion based representation learning. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 24963–24982. PMLR, 23–29 Jul 2023. URL <https://proceedings.mlr.press/v202/mittal23a.html>.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In Pierre Isabelle, Eugene Charniak, and Dekang Lin (eds.), *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pp. 311–318, Philadelphia, Pennsylvania, USA, July 2002. Association for Computational Linguistics. doi: 10.3115/1073083.1073135. URL <https://aclanthology.org/P02-1040>.
- Vaishnavi Patil, Matthew Evanusa, and Joseph JaJa. Dot-vae: Disentangling one factor at a time. In *International Conference on Artificial Neural Networks*, pp. 109–120. Springer, 2022.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International conference on machine learning*, pp. 1278–1286. PMLR, 2014.
- Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space. *Advances in Neural Information Processing Systems*, 34:11287–11302, 2021.
- Yiqi Wang, Wentao Chen, Xiaotian Han, Xudong Lin, Haiteng Zhao, Yongfei Liu, Bohan Zhai, Jianbo Yuan, Quanzeng You, and Hongxia Yang. Exploring the reasoning abilities of multimodal large language models (mllms): A comprehensive survey on emerging trends in multimodal reasoning. *arXiv preprint arXiv:2401.06805*, 2024.
- Jiayang Wu, Wensheng Gan, Zefeng Chen, Shicheng Wan, and S Yu Philip. Multimodal large language models: A survey. In *2023 IEEE International Conference on Big Data (BigData)*, pp. 2247–2256. IEEE, 2023.
- Shukang Yin, Chaoyou Fu, Sirui Zhao, Ke Li, Xing Sun, Tong Xu, and Enhong Chen. A survey on multimodal large language models. *arXiv preprint arXiv:2306.13549*, 2023.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. BERTscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*, 2019.

A EXAMPLES OF EXPLANATION GENERATION WITH CLEAR AND UNCLEAR PATTERNS

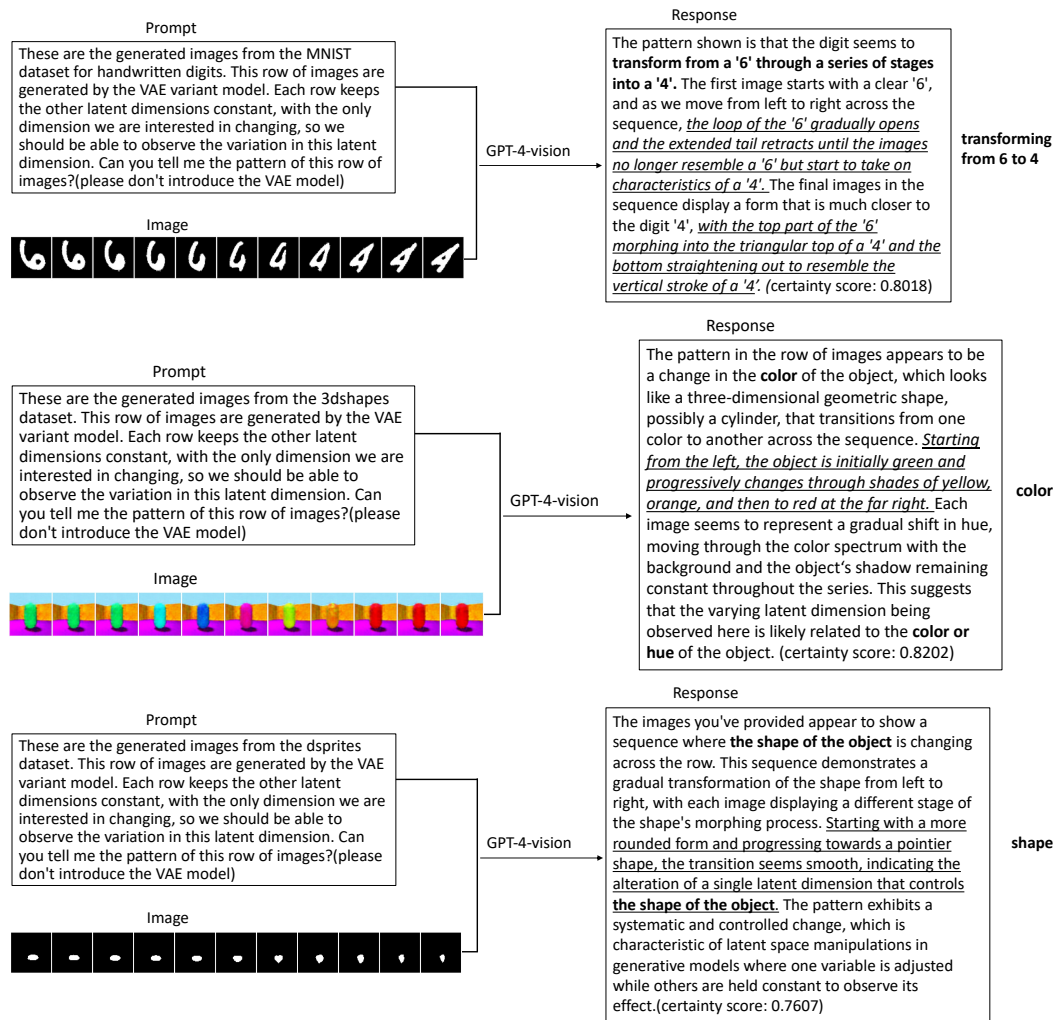


Figure 3: Sample images with clear patterns and sample prompts for GPT-4-vision to generate explanations

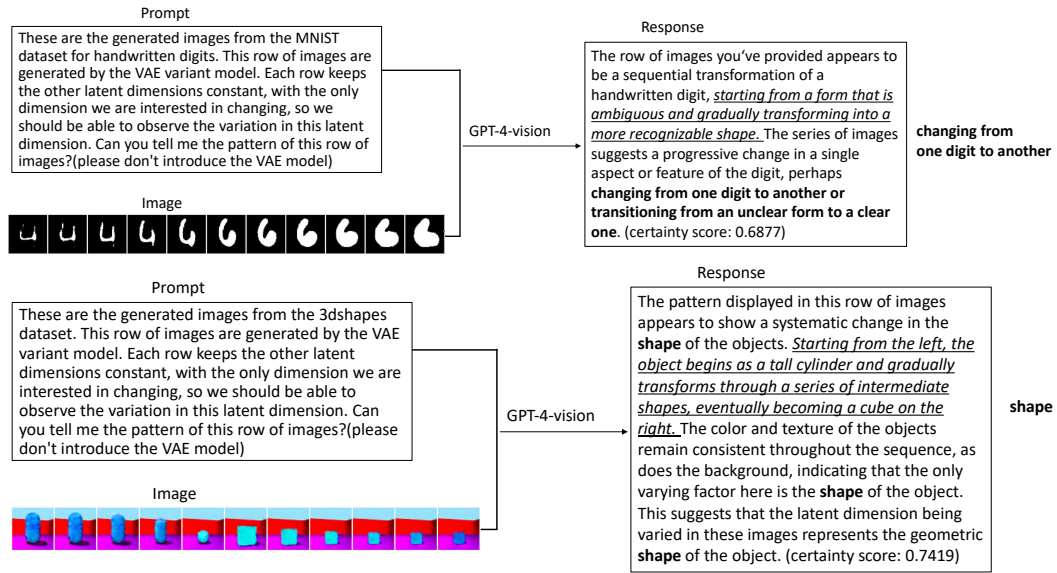


Figure 4: Sample images with unclear patterns and sample prompts for GPT-4-vision to generate explanations

B EXAMPLES OF EXPLANATIONS GENERATED BY DIFFERENT LMMs

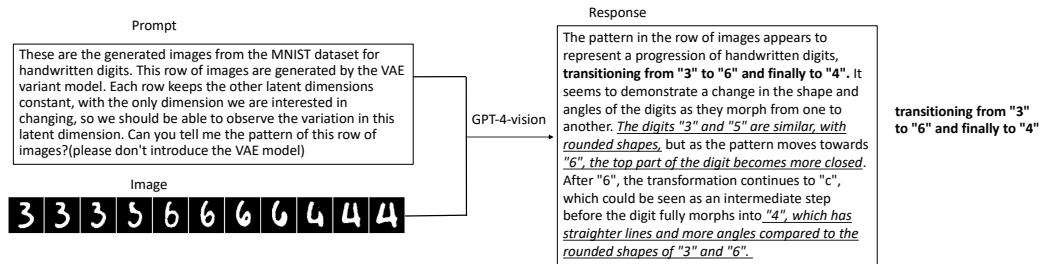


Figure 5: The sample explanation generated by GPT-4-vision for the MNIST dataset.

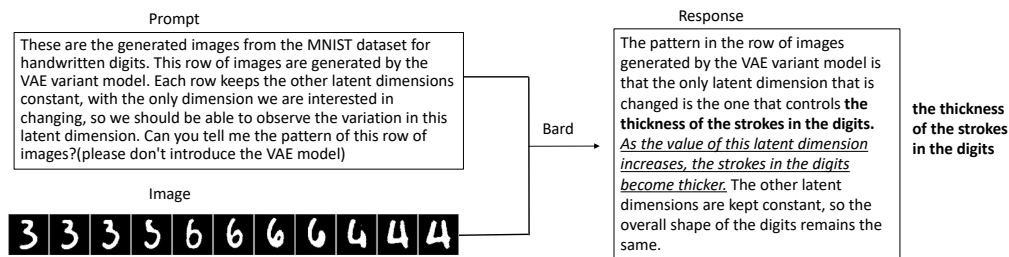


Figure 6: The sample explanation generated by Bard for the MNIST dataset.

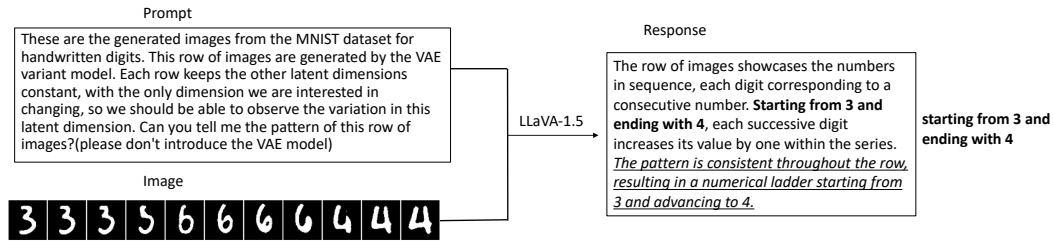


Figure 7: The sample explanation generated by LLaVA-1.5 for the MNIST dataset.

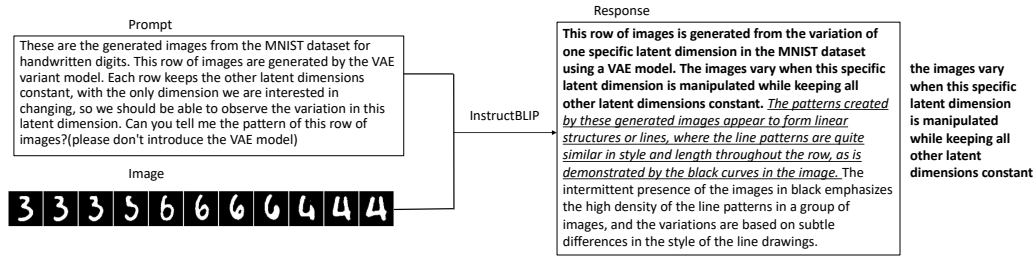


Figure 8: The sample explanation generated by InstructBLIP for the MNIST dataset.

C THE DISENTANGLEMENT EFFECTS OF DIFFERENT GENERATIVE MODELS ON EXPLANATIONS

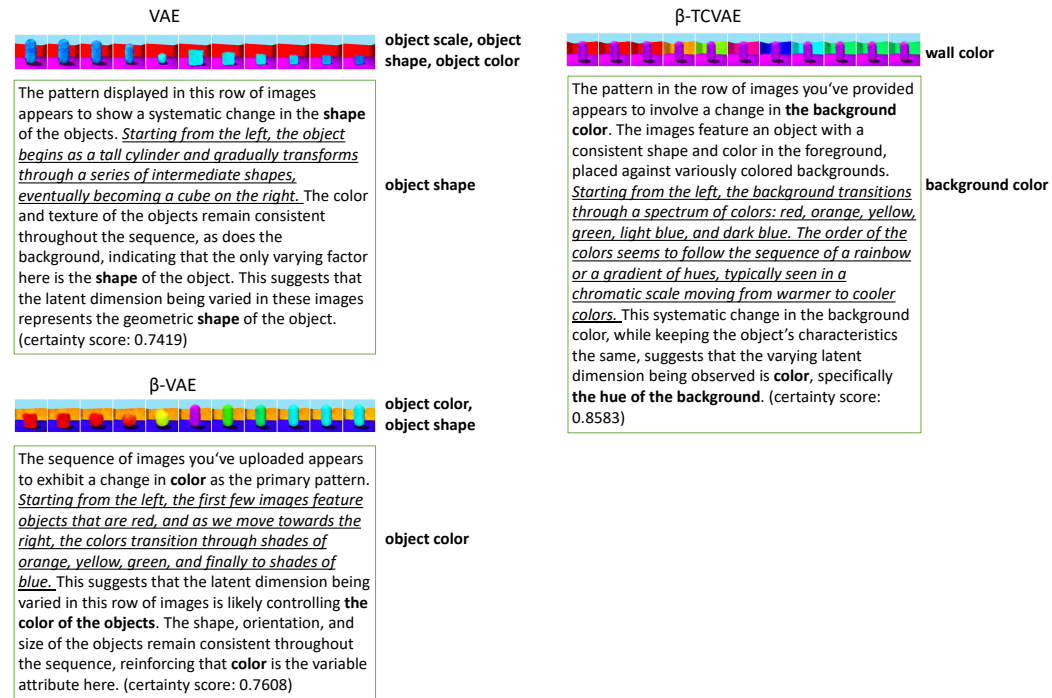


Figure 9: The impact of disentanglement in different generative models on explanations. The ground truth latent factors in the image sequences and the latent variables in the explanations are highlighted in **bold**, and the patterns of the latent variables are in *italics and underlined*.