Intra-Event Aware Imitation Game for Fast Detector Simulation

Hosein Hashemi
Faculty of Physics
Ludwig Maximilians University of Munich, Germany
Germany
gh.hashemi@physik.uni-muenchen.de

Nikolai Hartmann
Faculty of Physics
Ludwig Maximilians University of Munich, Germany
nikolai.hartmann@physik.uni-muenchen.de

Sahand Sharifzadeh
Faculty of Computer Science
Ludwig Maximilians University of Munich, Germany
sahand.sharifzadeh@gmail.com

James Kahn
Helmholtz AI
Karlsruhe Institute of Technology (KIT), Germany
james.kahn@kit.edu

Thomas Kuhr
Faculty of Physics
Ludwig Maximilians University of Munich, Germany
thomas.Kuhr@lmu.de

Abstract

While realistic detector simulations are an essential component of particle physics experiments, current methods are computationally inefficient, requiring significant resources to produce, store, and distribute simulation data. In this work, we propose the Intra-Event Aware GAN (IEA-GAN), a deep generative model which allows for faster and more resource-efficient simulations. We demonstrate its use in generating sensor-dependent images for the Pixel Vertex Detector (PXD) at the Belle II Experiment, the sub-detector with the highest spatial resolution. We show that using the domain-specific relational inductive bias introduced by our Relational Reasoning Module, one can approximate the concept of a collision event in the detector simulation. We also propose a Uniformity loss to maximize the information entropy of the IEA-GAN discriminator's knowledge and an Intra-Event Aware loss for the generator to imitate the discriminator's dyadic class-to-class knowledge. We show that the IEA-GAN not only captures the fine-grained semantic and statistical similarity between the images but also finds correlations among them, leading to a significant improvement in image fidelity and diversity compared to the previous state-of-the-art models.

Machine Learning and the Physical Sciences workshop, NeurIPS 2022.

1 Introduction

The "Fast Simulation" campaign in Particle Physics [1] sparked the search for faster and more storage-efficient simulation methods for collider physics experiments. These simulations play an essential role in various downstream tasks, including optimizing detector geometry, designing physics analyses, and searching for new phenomena beyond the Standard Model (SM). To simulate a sufficient amount of detector background at the Belle II [2] experiment, Generative Adversarial Networks (GANs) [3] have been explored for the first time [4, 5]. Since the introduction of class-conditional Generative Adversarial Networks [6], many schemes for capturing the class conditions in a more sophisticated way have been proposed. However, they are prone to class confusion [7] when they face a fine-grained dataset and fail to capture inter-class dependencies in a manner suitable for simulating multiple detector components within a single physics collision event.

In this work, we introduce Intra-Event Aware GAN (IEA-GAN), a novel method to capture intra-event dependencies and to generate correlated sensor-dependent Pixel Vertex Detector (PXD) [8] background images at Belle II with the highest fidelity while satisfying all relevant metrics. As a result, we propose new methods for training conditional GANs that not only seize both statistical-level and semantic-level information but also capture correlation among samples in a fine-grained image generation task.

2 Dataset

The Pixel Vertex Detector (PXD) [8], the innermost sub-detector of the Belle II experiment, measures the position of traversing charged particles originating from particle collision events by ionization of semi-conductor sensors in order to perform precise reconstruction of decay vertices. The configuration of the PXD consists of 40 modules within two layers, with 16 modules assembled in 8 planar ladders on the inner layer and 24 modules in 12 ladders in the outer. Each module consists of a pixel matrix sensor, resulting in a 250×768 hitmap image.

The PXD images are spatially asymmetric, have high resolution, and have statistical and semantic similarities. The occupancy of each image, defined as the fraction of non-zero valued pixels, shows statistical similarities between the two sensor layers and among the sensors in each layer. At the semantic and visual level, as the PXD images show extreme similarity, they can be classified as fine-grained images with small inter-class and significant intra-class variations [9], culminating in class confusion [7] because the discrimination between generated classes that are visually similar becomes very hard. Furthermore, since the information in an event comes from a single readout window for the PXD, the background processes happening in this window affect all sensors simultaneously, leading to a collective correlation among them as shown in fig. 3.

The data in each collision event in this work consists of 40 sensor-dependent grayscale 256×768 images[1]. For training, we used $40\,000$ simulated events, and for evaluation we used $10\,000$ simulated events. The dataset was produced with the basf2 [10] software, where the detector simulation was carried out using Geant4 [11]. For sampling, we consider the data sample as a single Monte-Carlo simulated collision event and each sensor as its class, such that in each sample, we have 40 unique classes of images corresponding to all 40 sensors.

3 Methods

3.1 Relational Reasoning Module

We propose the "Relational Reasoning Module" (RRM), a trainable layer that can capture contextualized embeddings from each image or class and is specifically designed to be compatible with GAN training policies. The outputs of this module are the updated features of each node, weighted by their importance or similarity. In the IEA-GAN Discriminator, this module takes the set of image embeddings from a sample as input nodes within a fully connected event graph, applies a pre-norm self-attention [12] over them, then updates each node's embedding via the attentive message

[1] Images are zero-padded on both sides from their original size of 250×768 in order to be divisible by 16, for training purposes.

propagation [13, 14, 15]. It then compactifies the information by projecting the normalized event graph onto a Hypersphere as a compact manifold.

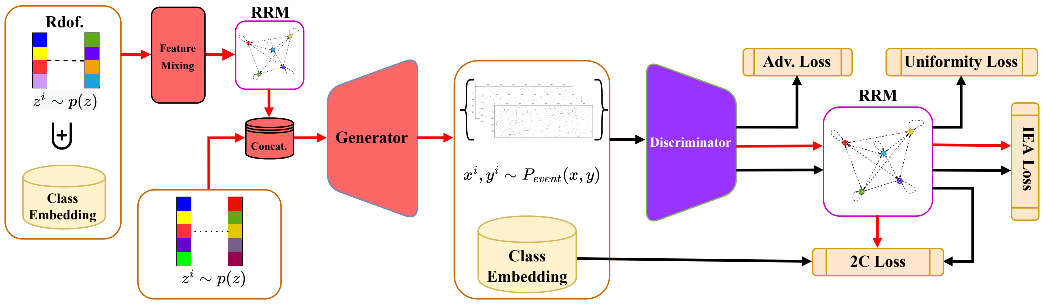

Figure 1: IEA-GAN architecture. The red lines correspond to the forward and backward passes of the Generator. The black lines correspond to the forward and backward passes of the discriminator. The discriminator is trained via the Adversarial Hinge loss [16] (Adv. Loss), 2C loss [17] and the Uniformity loss (see eq. (2)). On the other hand, the generator uses the Adversarial loss, 2C loss [17], and the IEA-loss (see eq. (1)).

In the IEA-GAN Generator, the RRM acts on all class token embeddings [2] and their learnable contextual embedding for each event. As conditions change for each event, having merely a class embedding, as is used in conditional GANs [6], is not enough. Thus, at each event, the Generator samples from a per-event shared normal distribution $r_i \sim \mathcal{N}(0, 1)$ as random degrees of freedom (Rdof). This acts as learnable contextual embedding for each class token that will be fused to it via the feature mixing layer.[3] Then, the IEA-GAN Generator does a message propagation amongst the fused tokens with the RRM. This way, we ensure that the Generator is aware of intra-event local changes. As a result, the Generator generates images that show intra-event correlations.

3.2 Intra-Event Aware Loss

To transfer the intra-event knowledge of the discriminator to the Generator in a more direct way, we introduce an Intra-Event Aware (IEA) loss for the Generator that captures class-to-class relations.

$$\ell_{\text{IEA}}(\mathbf{x}_r, \mathbf{x}_f) = \sum_{i,j} D_{\text{KL}}\left(\sigma\left(\mathbf{h}(x_i^{(r)})^\top \mathbf{h}(x_j^{(r)})\right) \middle\| \sigma\left(\mathbf{h}(x_i^{(f)})^\top \mathbf{h}(x_j^{(f)})\right)\right), \tag{1}$$

where $\mathbf{x}_r = \{x_i^{(r)}\}_{i=1}^m$ is the set of real images, and $\mathbf{x}_f = \{G(z^i, y^i, r^i) = x_i^{(f)}\}_{i=1}^m$ the set of generated images. The softmax function, $\sigma : \mathbb{R}^m \to [0, 1]^m$, normalizes the dot-product self-attention between the image embeddings. The map $\mathbf{h} : \mathbb{R}^{k\times m} \to \mathbb{S}^{n\times m}$ is the unit hypersphere projection of the discriminator. $D_{\text{KL}}(.||.)$ is the KL divergence [18] which takes two $m \times m$ matrices that have values in the closed unit interval (due to the softmax function). By minimizing ℓ_{IEA}, we are putting a discriminator-supervised penalizing system over the intra-event awareness of the Generator, encouraging it to look for more detailed dyadic connections [19] among the images. Ultimately, we want to maximize the consensus between the hyperspheres of real and generated image embeddings. The Intra-Event Aware loss will be added to the Generator's losses (The Adversarial Hinge loss [16] and 2C Loss [17]).

3.3 Uniformity Loss

The other crucial loss function comes from contrastive representation learning. The task of learning fine-grained class-to-class relations among images means we also want to ensure the feature vectors have as much hyperspherical diversity as possible. Thus, by imposing a uniformity condition over the feature vectors on the unit hypersphere, they preserve as much information as possible since the uniform distribution carries a high entropy. This idea stems from the Thomson problem [20], where a

[2] One-Hot encoding of each PXD sensor number

[3] The feature mixing layer is a single linear layer without any activation functions.

static equilibrium with minimal potential energy is sought to distribute N electrons on the unit sphere in the evenest manner. To do this, we use the uniformity metric proposed by [21, 22], which is based on a Gaussian potential kernel,

$$\mathcal{L}_{\text{uniform}}(x;t) = \log \mathbb{E}_{x_i,x_j \sim p_{\text{data}}}[\exp(-t\|\mathbf{h}(x_i) - \mathbf{h}(x_j)\|_2^2)]. \tag{2}$$

Upon minimizing this loss for the discriminator, it tries to maintain a uniform distance between samples that are not well-clustered and thus not similar. This loss is necessary for capturing the exact distribution of the mean occupancy and balancing the inter-class pulling force of the RRM. The Uniformity loss will be added to the Discriminator's losses (The Adversarial Hinge loss [16] and 2C Loss [17]).

4 Result

We compare IEA-GAN with four state-of-the-art conditional image generation baselines and the Geant4-simulated reference: BigGAN-deep [16], ContraGAN [17], and our previous works, PE-GAN [5] and WGAN-gp [4][4]. As we are interested in optimal pixel-level properties, diversity, and correlations of generated images while minimizing generator complexity due to computational limitations, we choose the model's weight with the best Frechet Inception Score (FID) [23]. Based on the recent Clean-FID project [24], we fine-tune the Inception-V3 [25] model on the PXD images, as the PXD images are very different from the natural images used in its initial training. Table 1 shows that IEA-GAN images have the lowest FID score, ergo better diversity and fidelity compared to the other models. Furthermore, since in GAN training there is no meaningful way to define a minimal loss, our stopping point is the divergence of the FID which is very correlated with the quality of other metrics.

	WGAN-gp	PE-GAN	BigGAN-deep	ContraGAN	IEA-GAN
FID	12.09	5.21	4.40	2.54 ± 0.43	$\mathbf{1.50 \pm 0.16}$

Table 1: FID comparison between models (all models in the benchmark are highly tuned to the current problem and dataset). The lower the FID, the better the image quality and diversity.

At the image level, we look at the Pixel Intensity Distribution, the Occupancy Distribution, and the Mean Occupancy. The Pixel Intensity Distribution is the cumulative sum of pixel values over the given number of events which defines the distribution of the background hit energy. The Occupancy Distribution and the Pixel Intensity Distribution are evaluated over all sensors of a given number of events, while the Mean Occupancy corresponds to the mean value of sparsity across events for each sensor. This pixel-level information defines each sensor's sparsity, hence background levels, during physics analysis with the basf2 software [10] at Belle II. The Pixel Intensity Distribution and Occupancy Distribution are shown in fig. 2.

The aforementioned metrics are equivariant under permutation between the samples among events. Therefore, we compute Spearman's correlation between the sensors' occupancy along the population of generated events. Although the desired correlation is different from IEA-GAN, as shown in fig. 3, IEA-GAN understands a positive correlation for intra-layer sensors and a primarily negative correlation for inter-layer sensors. At the Physics-reconstruction level with the basf2 software [10], the discrepancy in the correlation affects the correspondence between the underlying Background processes that are responsible for the PXD hits.

5 Conclusions and Outlook

In this work, we have proposed a novel way to generate correlated high-resolution particle detector background data. IEA-GAN captures the pair-wise class-to-class relations and exhibits intra-event correlations among the generated high-resolution images in an end-to-end way for a fine-grained conditional image generation task. In the future, we think that injecting the physics information and mask-conditioning over them while respecting the symmetries of the new modality would solve

[4]WGAN-gp are used only for FID

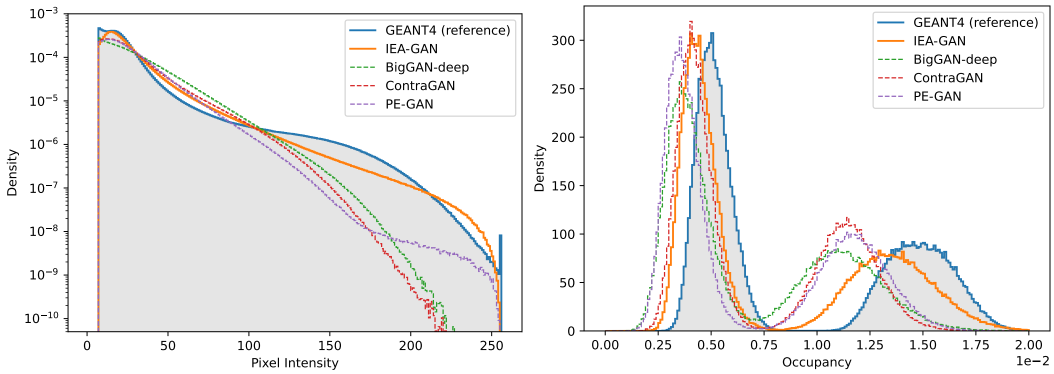

Figure 2: Pixel Intensity Distribution in log scale (left) and Distribution of the occupancy for 10 000 events

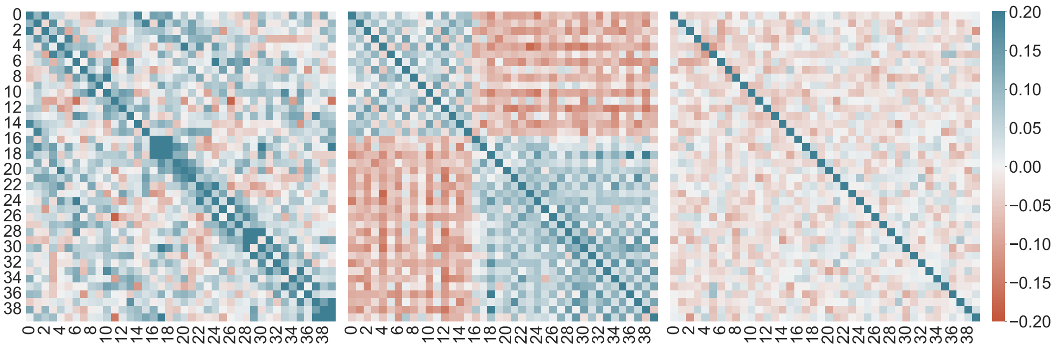

Figure 3: Spearman's correlation between the occupancy of real sensor images (left), sensor images from IEA-GAN (center), and sensor images from ContraGAN (right).

the remaining discrepancy between the real and generated images. As the fine-grained conditional generation of collider events is omnipresent in all fast and efficient detector simulations, we believe that the Intra-Event Aware GAN (IEA-GAN) offers a robust solution for other particle physics experiments and simulations at Belle II [2] and LHC [26].

References

[1] Sofia Vallecorsa. Generative models for fast simulation. *Journal of Physics: Conference Series*, 1085:022005, September 2018. 2

[2] Tetsuo Abe, I Adachi, K Adamczyk, S Ahn, H Aihara, K Akai, M Aloi, L Andricek, K Aoki, Y Arai, et al. Belle ii technical design report. *arXiv preprint arXiv:1011.0352*, 2010. 2, 5

[3] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Networks, June 2014. 2

[4] Matej Srebre, Pascal Schmolz, Hosein Hashemi, Martin Ritter, and Thomas Kuhr. Generation of Belle II Pixel Detector Background Data with a GAN. *EPJ Web Conf.*, 245:02010, 2020. 2, 4

[5] Hosein Hashemi, Nikolai Hartmann, Thomas Kuhr, Martin Ritter, and Matej Srebre. Pixel Detector Background Generation using Generative Adversarial Networks at Belle II. *EPJ Web of Conferences*, 251:03031, January 2021. 2, 4

[6] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014. 2, 3

[7] Minguk Kang, Woohyeon Shim, Minsu Cho, and Jaesik Park. Rebooting ACGAN: Auxiliary Classifier GANs with Stable Training. In *Advances in Neural Information Processing Systems*, volume 34, pages 23505–23518, Virtual, 2021. Curran Associates, Inc. 2

[8] F. Mueller. Some aspects of the Pixel Vertex Detector (PXD) at Belle II. *Journal of Instrumentation*, 9(10):C10007–C10007, October 2014. 2

[9] Xiu-Shen Wei, Yi-Zhe Song, Oisin Aodha, Jianxin Wu, Yuxin Peng, Jinhui Tang, Jian Yang, and Serge Belongie. Fine-Grained Image Analysis with Deep Learning: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PP:1–1, November 2021. 2

[10] Nils Braun, T. Kuhr, M. Ritter, C. Pulvermacher, and Thomas Hauth. The Belle II Core Software. *Computing and Software for Big Science*, 3, November 2018. 2, 4

[11] Sea Agostinelli, John Allison, K al Amako, John Apostolakis, H Araujo, Pedro Arce, Makoto Asai, D Axen, Swagato Banerjee, GJNI Barrand, et al. Geant4—a simulation toolkit. *Nuclear instruments and methods in physics research section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250–303, 2003. 2

[12] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, pages 6000–6010, Red Hook, NY, USA, December 2017. Curran Associates Inc. 2

[13] Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018. 3

[14] Sahand Sharifzadeh, Sina Moayed Baharlou, and Volker Tresp. Classification by Attention: Scene Graph Classification with Prior Knowledge. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(6):5025–5033, May 2021. 3

[15] Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-centric learning with slot attention. *Advances in Neural Information Processing Systems*, 33:11525–11538, 2020. 3

[16] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2019. 3, 4

[17] Minguk Kang and Jaesik Park. ContraGAN: Contrastive Learning for Conditional Image Generation. In *Advances in Neural Information Processing Systems*, volume 33, pages 21357–21369, Virtual, 2020. Curran Associates, Inc. 3, 4

[18] S. Kullback and R. A. Leibler. On Information and Sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, March 1951. 3

[19] Longbing Cao. Coupling learning of complex interactions. *Information Processing & Management*, 51(2):167–186, March 2015. 3

[20] J.J. Thomson. XXIV. On the structure of the atom: An investigation of the stability and periods of oscillation of a number of corpuscles arranged at equal intervals around the circumference of a circle; with application of the results to the theory of atomic structure. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 7(39):237–265, March 1904. 3

[21] Tongzhou Wang and Phillip Isola. Understanding Contrastive Representation Learning through Alignment and Uniformity on the Hypersphere. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 9929–9939, Vienna, Austria, November 2020. PMLR. 4

[22] A. Kuijlaars and E. Saff. Asymptotics for minimal discrete energy on the sphere. *Transactions of the American Mathematical Society*, 350(2):523–538, 1998. 4

[23] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, pages 6629–6640, Red Hook, NY, USA, December 2017. Curran Associates Inc. 4

[24] Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On Aliased Resizing and Surprising Subtleties in GAN Evaluation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11410–11420, 2022. 4

[25] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the Inception Architecture for Computer Vision. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826, June 2016. 4

[26] Lyndon Evans and Philip Bryant. LHC Machine. *Journal of Instrumentation*, 3(08):S08001–S08001, August 2008. 5

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [N/A]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No] We don't publish the code along with this contribution, but will do so in the future.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [No]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See table 1
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] The dataset was created using the Belle II software basf2, which is referenced.
   (b) Did you mention the license of the assets? [No] We don't publish the dataset along with this contribution, but may do so in the future. The basf2 software used to create the dataset is licensed under LGPLv3.
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [No] But the authors are members of the Belle II collaboration and are as such allowed to use the data.
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]