Short Paper

Generating Artificial Patients With Reliable Clinical Characteristics Using a Geometry-Based Variational Autoencoder: Proof-of-Concept Feasibility Study

Fabrice Ferré¹, MD, PhD; Stéphanie Allassonnière², PhD; Clément Chadebec², PhD; Vincent Minville¹, MD, PhD

¹Department of Anesthesia, Intensive Care and Perioperative Medicine, Purpan University Hospital, Toulouse, France

²Université Paris Cité, Unité Mixte de Recherche S1138, Institut national de recherche en sciences et technologies du numérique, Sorbonne University, Paris, France

Corresponding Author:

Fabrice Ferré, MD, PhD Department of Anesthesia, Intensive Care and Perioperative Medicine Purpan University Hospital Place du Dr Baylac Toulouse, 31300 France Phone: 33 561779988 Email: <u>fabriceferre31@gmail.com</u>

Abstract

Background: Artificial patient technology could transform health care by accelerating diagnosis, treatment, and mapping clinical pathways. Deep learning methods for generating artificial data in health care include data augmentation by variational autoencoders (VAE) technology.

Objective: We aimed to test the feasibility of generating artificial patients with reliable clinical characteristics by using a geometry-based VAE applied, for the first time, on high-dimension, low-sample-size tabular data.

Methods: Clinical tabular data were extracted from 521 real patients of the "MAX" digital conversational agent (BOTdesign) created for preparing patients for anesthesia. A 3-stage methodological approach was implemented to generate up to 10,000 artificial patients: training the model and generating artificial data, assessing the consistency and confidentiality of artificial data, and validating the plausibility of the newly created artificial patients.

Results: We demonstrated the feasibility of applying the VAE technique to tabular data to generate large artificial patient cohorts with high consistency (fidelity scores>94%). Moreover, artificial patients could not be matched with real patients (filter similarity scores>99%, κ coefficients of agreement<0.2), thus guaranteeing the essential ethical concern of confidentiality.

Conclusions: This proof-of-concept study has demonstrated our ability to augment real tabular data to generate artificial patients. These promising results make it possible to envisage in silico trials carried out on large cohorts of artificial patients, thereby overcoming the pitfalls usually encountered in in vivo trials. Further studies integrating longitudinal dynamics are needed to map patient trajectories.

(J Med Internet Res 2025;27:e63130) doi: 10.2196/63130

KEYWORDS

digital health; artificial data; variational autoencoder; data science; artificial intelligence; health monitoring; deep learning; medical imaging; imaging; magnetic resonance imaging; Alzheimer disease; anesthesia; prediction; data augmentation

Introduction

RenderX

With the growing impact of data science technologies, novel health care ecosystems centered around artificial patients are developing. The data science–based approach for generating artificial patients involves augmenting real data [1]. Thus, new

https://www.jmir.org/2025/1/e63130

artificial data are created with characteristics similar to those of the original population of interest. Such data could be particularly valuable in clinical research, offering the potential for studies that are not only more cost-effective but, more importantly, also more inclusive and impactful—especially in cases where patient recruitment poses a challenge [2].

From a methodological point of view, artificial labeled data are generated using mechanistic or statistical methods. The mechanistic approach combines known equations from physical, biological, or other fields to describe a phenomenon, referred to as digital twins. This approach could facilitate personalized therapeutics [3], though its routine use in medical processes remains limited [4]. Deep learning methods for generating artificial data in health care include technologies based on generative adversarial networks (GANs) [5] or variational autoencoders (VAEs) [1,6]. However, most of the studies using GANs focused on a fairly large training set (over 1000 training samples) or on low-dimension data, while it remains very challenging to gather such large cohorts of labeled patients. Therefore, the case of high-dimensional data combined with a small sample size (a situation commonly encountered in medicine) remains largely unexplored by GAN technologies. Chadebec et al [1] have recently demonstrated, by using a VAE, that the artificial augmentation of medical imaging data significantly improved classification accuracy. The balanced accuracy increases from 66% to 74% for a convolutional neural network classifier trained with small datasets (50 magnetic resonance images each of cognitively healthy individuals and patients with Alzheimer disease), while improving greatly the sensitivity and specificity of the classification metrics [1]. In other words, a geometry-based VAE was able to produce meaningful samples from high-dimension, low-sample-size (HDLSS) imaging datasets. This method, validated for an image classification task (voxels), deserves to be tested on (clinical) tabular data. Given the multiplicity and complexity of the data obtained in anesthesia, this statistical approach using artificially augmented data could be of major interest to identify predictive clinical factors of poor outcomes with accuracy and reliability. Thus, the aim of our study was to test the feasibility of generating artificial patients with reliable clinical characteristics by using a VAE applied on HDLSS tabular data.

Methods

Overview

Clinical tabular data were extracted from 521 real patients of the MAX (BOTdesign) database. MAX is a digital conversational agent for preparing patients for anesthesia [7,8]. Collected data included demographic characteristics, past medical history, medication, and other relevant medical items.

Details of the methodological approach used to train the model and generate artificial data are available in Multimedia Appendix 1. Briefly, the dataset included 521 patients about to undergo anesthesia, each with 85 clinical features. Once data preprocessing was completed (Figure 1), the model was trained on a (521, 103) dimension dataset using a modified Pyraug's training pipeline (Figure 2). Training hyperparameters were set to 1000 epochs, a batch size of 32, and a learning rate of 0.001. Two datasets of 5000 and 10,000 artificial patients were generated, representing a data increase rate of 10 and 20 artificial patients, respectively, for 1 real patient.

The next step involved assessing the consistency (fidelity scores) and confidentiality (filter similarity scores and degree of anonymization) of artificial data.

Finally, the plausibility of newly generated artificial patients was validated through expert human supervision (categorization task).

ood type AB+

0

1

0

Gender	Height	МЕТ	Blood type	
М	175	< 4	0+	
F	165	4 - 7	AB+	
F	155	> 7	A-	
Binarization	Normalization	Ordinal encodin + Normalization	g One-hot encoding	
Gender	Height	MET	Blood type O+	B
0	1	0	1	
1	0.5	0.5	0	
1	0	1	0	

Figure 1. Illustration of the data preprocessing step. MET: metabolic equivalent task.



Blood type A-

0

0 1

Figure 2. Illustration of the complete pipeline for generating artificial patients. For more details, see Multimedia Appendix 1. TDProcessor: TabularDataProcessor.



Ethical Considerations

We did not seek ethics approval in accordance with the Commision Nationale Informatique et Libertés policy on secondary analyses of preexisting datasets (Titles I and II) [9], as patients were informed that their data may be securely stored, coded for confidentiality, and used for research unless they explicitly object.

Results

Consistency and Confidentiality of Artificial Data

To assess the consistency of newly generated artificial data, fidelity scores were calculated. A fidelity score is defined as the arithmetic mean of its 3 components: numerical and categorical data distribution stability and numerical data correlation stability. Results are presented in Table 1 and illustrated in Figures 3 and 4. Equations used to calculate stability scores for numerical and categorical data are available in Multimedia Appendix 1.

Table 1. Fidelity scores and their components for the 5000 and 10,000 artificial patients generated.

Scores	5000 artificial patients, %	10,000 artificial patients, %
Fidelity score	97.8	94.6
Numerical stability	100	100
Categorical stability	96.4	91.2
Numerical correlation stability	97.1	92.6



Figure 3. Distribution of (A) the numerical variable "age" and (B) the categorical variable "gender" (in relative percentage) from the datasets of 521 real and 5000 artificial patients.



Figure 4. Matrices illustrating the correlation of numerical variables from the datasets of 521 real and 5000 artificial patients (A) and (B) and their differences (C).



With fidelity scores of 97.8% and 94.6%, data of the 5000 and 10,000 artificial patients were considered realistic and representative of the real data.

To assess the confidentiality of newly generated artificial data, filter similarity scores (proportion of data not similar to the initial dataset) and degree of anonymization (Euclidean distance) were calculated. The results strongly confirm the nonsimilarity of the artificial data with the initial real data (filter similarity scores>99.9%) and a high degree of anonymization (the artificial data were further away from the initial dataset than the initial data were from each other).

Plausibility of Artificial Patients

A categorization task was performed by 3 experienced anesthetists who were blinded, with the distribution of a balanced sample of 100 real and artificial patients generated using the VAE. Anesthetists were asked to determine whether each patient was real or artificial. The κ coefficients of agreement were –0.12 (95% CI –0.31 to 0.07), 0.15 (95% CI 0.1-0.26), and 0.09 (95% CI –0.14 to 0.15). Given the very low agreement coefficients (<0.2), none of the 3 experts could differentiate between real and artificial patients, arguing for the medical plausibility of the artificial patients.

Discussion

Through this proof-of-concept study, we demonstrated, for the first time, the feasibility to transpose the VAE technique from imaging to HDLSS clinical data for the generation of a large number of artificial patients. The high fidelity scores obtained demonstrate the consistency of our artificial cohorts. Moreover, as suggested by the high filter similarity scores and the low agreement coefficients of the categorization task, the artificial patients could not be matched with real patients.

The use of artificial intelligence (AI) in health care presents particularly complex challenges inherent in the types of data it relies on (sensitive, sparse, heterogeneous, limited, etc). However, research efforts over the last few years, which have focused on meeting these specific challenges, now offer a glimpse of dizzying potential, particularly in the field of clinical research. In silico trials on artificial patient cohorts can quickly and cost-effectively include diverse minority groups (eg, rare

RenderX

diseases, children, pregnant women, and ethnic minorities), reducing risks and recruitment challenges of traditional in vivo trials. In this setting, artificial patients' technology has the potential to transform health care by improving diagnosis, treatment, and mapping clinical pathways [3].

Validation of mathematical models and algorithms is important from an ethical point of view. Like any data processing, the implementation of artificial patient cohorts requires human oversight (in reference to the AI Act [10]). This principle of "Human Guarantee" refers to the need not to relinquish decision-making autonomy in the context of increasingly rapid dissemination of AI. Organized in the form of human oversight committees, these control measures will help better understand the phase of modeling artificial patient populations and ensure that it is as unbiased and reliable as possible. The need to protect health data and the rights of the individuals is a matter of debate that will likely be resolved with the expected advances in anonymization. Ensuring that artificial and real data are not similar is one way of guaranteeing anonymization and confidentiality. In this setting, no real patient should be identifiable from artificial data.

Our results must be interpreted with caution and a number of limitations should be borne in mind. First, our study presents promising results but does not present a comparative analysis with existing methods for generating artificial patient data (eg, GANs). In this setting, further studies designed to identify the most effective approach could be of major interest. Second, our study strongly lacks a comparative analysis with the late postoperative outcomes of real patients. Further studies integrating longitudinal dynamics are needed to map trends and identify patient trajectories. In this context, the recent update of MAX with the implementation of a postoperative digital conversational agent for the collection of recovery data based on patient-reported outcome measures could be of major interest. Indeed, we plan to compare longitudinal data from real patients with those obtained in a cohort of artificial patients. Finally, the results of our proof-of-concept study are encouraging, but it seems necessary to replicate the method. For instance, a process is underway to validate the reliability of artificial patient cohorts by replaying old clinical trials using data shared on the Yale University Open Data Access. Moreover, we plan to apply our VAE method to other kinds of data (eg, multimodal data). In

this setting, the first European web platform for augmenting data and creating artificial patient cohorts with VAE's generative AI has recently been created (ORIGA; BOTdesign).

To conclude, we demonstrate, for the first time, the feasibility to augment HDLSS clinical tabular data by using a VAE. The newly generated artificial patient cohorts were consistent with real source data. We believe that in silico trials can be used to track a variety of health indicators and generate key insights. Artificial patients will revolutionize health care paving the way for a more precise, personalized, and predictive medicine.

Acknowledgments

We would like to thank Open Studio (IA-Lab, Lyon, France) for their data analyses. We would also like to thank Olivier Thuillart and Jean-Louis Fraysse from BOTdesign (Toulouse, France), the company that created the digital companion MAX. The research leading to these results has received funding from the French government under management of Agence Nationale de la Recherche as part of the "Investissements d'avenir" program (reference ANR-19-P3IA-0001, PRAIRIE 3IA Institute; and reference ANR-10-IAIHU-06, Agence Nationale de la Recherche-10-IA Institut Hospitalo-Universitaire-6).

Data Availability

Data are available from the corresponding author upon request.

Authors' Contributions

Conceptualization: FF, SA, and VM Formal analysis: SA and CC Investigation: FF and VM Methodology: FF and SA Writing—original draft: FF, SA, CC, and VM Writing—review and editing: FF

Conflicts of Interest

None declared.

Multimedia Appendix 1

Additional material. [DOCX File , 176 KB-Multimedia Appendix 1]

References

- Chadebec C, Thibeau-Sutre E, Burgos N, Allassonniere S. Data augmentation in high dimensional low sample size setting using a geometry-based variational autoencoder. IEEE Trans Pattern Anal Mach Intell. Mar 2023;45(3):2879-2896. [doi: 10.1109/TPAMI.2022.3185773] [Medline: 35749321]
- 2. Desai M. Recruitment and retention of participants in clinical studies: critical issues and challenges. Perspect Clin Res. 2020;11(2):51. [doi: 10.4103/picr.picr_6_20]
- 3. Elkefi S, Asan O. Digital twins for managing health care systems: rapid literature review. J Med Internet Res. Aug 16, 2022;24(8):e37641. [FREE Full text] [doi: 10.2196/37641] [Medline: 35972776]
- 4. Fischer R, Volpert A, Antonino P, Ahrens TD. Digital patient twins for personalized therapeutics and pharmaceutical manufacturing. Front Digit Health. 2023;5:1302338. [FREE Full text] [doi: 10.3389/fdgth.2023.1302338] [Medline: 38250053]
- Segal B, Rubin DM, Rubin G, Pantanowitz A. Evaluating the clinical realism of synthetic chest x-rays generated using progressively growing GANs. SN Comput Sci. 2021;2(4):321. [FREE Full text] [doi: 10.1007/s42979-021-00720-7] [Medline: 34104898]
- 6. Allassonière S, Fraysse J. Artificial data in healthcare: analysis and areas for consideration [Article in French]. 2024. URL: https://static.botdesign.net/docs/VF Livre blanc Données de santé artificielles-250424.pdf [accessed 2025-03-19]
- Ferré F, Boeschlin N, Bastiani B, Castel A, Ferrier A, Bosch L, et al. Improving provision of preanesthetic information through use of the digital conversational agent "MyAnesth": prospective observational trial. J Med Internet Res. Dec 04, 2020;22(12):e20455. [FREE Full text] [doi: 10.2196/20455] [Medline: 33275108]
- 8. Ferré F, Laurent R, Furelau P, Doumard E, Ferrier A, Bosch L, et al. Perioperative risk assessment of patients using the MyRISK digital score completed before the preanesthetic consultation: prospective observational study. JMIR Perioper Med. Jan 16, 2023;6:e39044. [FREE Full text] [doi: 10.2196/39044] [Medline: 36645704]
- 9. Deliberation No. 2018-155 of May 3, 2018 approving the reference methodology relating to the processing of personal data implemented in the context of research not involving human beings, studies and evaluations in the field of health (MR-004). Commision Nationale Informatique et Libertés. URL: https://tinyurl.com/3c598h7y [accessed 2025-04-11]

RenderX

10. Article 14: Human Oversight. EU Artificial Intelligence Act. URL: <u>https://artificialintelligenceact.eu/article/14/</u> [accessed 2024-06-13]

Abbreviations

AI: artificial intelligence GAN: generative adversarial network HDLSS: high-dimension, low-sample-size VAE: variational autoencoder

Edited by A Coristine; submitted 11.06.24; peer-reviewed by D Chen, V Astha; comments to author 08.07.24; revised version received 01.11.24; accepted 17.12.24; published 17.04.25

<u>Please cite as:</u> Ferré F, Allassonnière S, Chadebec C, Minville V Generating Artificial Patients With Reliable Clinical Characteristics Using a Geometry-Based Variational Autoencoder: Proof-of-Concept Feasibility Study J Med Internet Res 2025;27:e63130 URL: <u>https://www.jmir.org/2025/1/e63130</u> doi: <u>10.2196/63130</u> PMID:

©Fabrice Ferré, Stéphanie Allassonnière, Clément Chadebec, Vincent Minville. Originally published in the Journal of Medical Internet Research (https://www.jmir.org), 17.04.2025. This is an open-access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work, first published in the Journal of Medical Internet Research (ISSN 1438-8871), is properly cited. The complete bibliographic information, a link to the original publication on https://www.jmir.org/, as well as this copyright and license information must be included.

