IDADA: A Blended Inductive-Deductive Approach for Data Augmentation

Pierangela Bruno[1][0000−0002−0832−0151], Francesco Calimeri[1][0000−0002−0866−0834], Francesca Filice[1], Cinzia Marte[1][0000−0003−3920−8186], and Simona Perri[1][0000−0002−8036−5709]

Department of Mathematics and Computer Science, Via P. Bucci, Rende, 87036, Italy
{pierangela.bruno,francesco.calimeri,cinzia.marte,simona.perri}@unical.it, filicefrancesca.ff@gmail.com

Abstract. This work proposes a hybrid approach to Data Augmentation that blends inductive and deductive reasoning. In particular, the approach effectively utilizes a modest collection of labeled images while employing logic programs to declaratively define the structure of new images, allowing for flexible and dynamic image generation; the use of logic programming ensures adherence to both domain-specific constraints and given desiderata. The resulting structures are then used for guiding the generation of new realistic images based on a dedicated Deep-Learning process. The general approach can be particularly of use in biomedical and healthcare scenarios, where building extensive datasets of quality images is in general a hard prerequisite for many applications that is challenging to meet. The approach is specialized to two real-world case studies featuring laryngeal endoscopic and cataract images, respectively, and experiments conducted for assessing the method are discussed.

Keywords: Data Augmentation · Hybrid Approaches · Deep Learning · Deductive Reasoning · Inductive Reasoning.

1 Introduction

In recent years, Deep Learning (DL) applications have gained significant attention for their impressive results in various fields such as image processing, pattern recognition, object recognition [32,27]. However, these methodologies depend on models that require training on proper background knowledge, which must be represented in datasets that are adequate in terms of size, quality, and various other factors. Obtaining a substantial amount of "good" training data can be challenging in certain domains; this is particularly common in biomedicine, due to factors such as accessibility, costs, manual annotation effort, data availability, and class imbalance. To address this challenge, data augmentation techniques have been extensively researched to enrich and enhance poor datasets. Generative models like Generative Adversarial Networks (GANs) have been proposed, in particular to create synthetic yet realistic images, showing significant potential. However, these approaches also suffer from drawbacks and limitations. For instance, their training can be unstable and slow [16]; moreover, guiding feature extraction and image generation typically relies on the composition and adaptation of the training dataset, making it challenging to leverage available knowledge and express preferences for data generation. Nonetheless, such knowledge can be valuable in avoiding the generation of erroneous images, reducing generation times, and enhancing overall result quality by ensuring reliable generation of images aligned with specific criteria.

In this work, we present IDADA, a framework blending inductive and deductive strategies for data augmentation. IDADA relies on Answer Set Programming (ASP), a purely declarative formalism rooted in logic programming and non-monotonic reasoning [4,17,18]. ASP offers explainability by design, suitability for complex Knowledge Representation and Reasoning (KRR) tasks, and benefits from available efficient implementations [22]; its use for data augmentation facilitates the expression of constraints arising from background knowledge and desired features, thus allowing for guiding the automatic generation of new data that comply with domain knowledge.

Currently, IDADA focuses on image data generation. Basically, starting from a set of real (or realistic) images, IDADA: (*i*) identifies (at a semantic level) a set of distinctive elements that are supposed to be present in the images to be created; (*ii*) produces a knowledge base by encoding in ASP needed domain knowledge along with constraints and desiderata about how the new images should appear, in terms of the identified elements; (*iii*) generates a number of image structures by placing instances of distinctive elements and arranging them according to the encoded knowledge base, thus obtaining new labeled images (that are, essentially, semantically segmented images); (*iv*) employs DL methods to "fill" the image areas, thus producing plausible synthetic images based on the labeled image structures already generated.

To the best of our knowledge, this work represents one of the pioneering attempts in using Answer Set Programming (ASP) for medical image generation and augmentation.

We experiment with IDADA for generating synthetic images in two biomedical scenarios: cataract images starting from the Cataract Dataset (CaDIS) [19] and laryngeal endoscopic images, starting from the Laryngeal Endoscopic Dataset [25]. The results are encouraging, and demonstrate the effectiveness of IDADA in enabling the generation of new images based on declaratively expressed criteria.

The remainder of the paper is structured as described next. In Section 2 we provide some background and discuss related works, before introducing IDADA, the herein proposed framework for data augmentation, in Section 3. In Section 4 we report on the result of the experimental campaign, and eventually present our conclusion and briefly discuss future perspectives in Section 5.

2 Related Work

Image data augmentation techniques have been widely studied in the literature and used in state-of-the-art solutions to reduce overfitting, increase generalizability, and overcome the lack of data or other limitations that could affect algorithm performance. Indeed, data augmentation: (*i*) results in general much less expensive than regular data collection with its label annotation; (*ii*) can be extremely accurate (it is generated from ground-truth data); (*iii*) is controllable, to some extent, in generating balanced data [23].

Typically, image data augmentation is performed relying on "classical" strategies or methods based on DL techniques. In the first case, geometric transformation (i.e., flipping, rotation, shearing, cropping, translation in the geometric transformation) and photometric shifting (i.e., color space shifting, image filtering, addition of noise) are applied to existing available images in order to enrich the collection [23]. However, these techniques present some disadvantages, including memory consumption, transformation costs, and additional training time. Also, some strategies, such as photometric shifting, can produce the eliminations of important color information or specific features in the image, thus not always guaranteeing the preservation of nature and meaning of the image labels [29]. Hence, DL-based methods have been increasingly employed. Indeed, DL methods, especially Generative Adversarial Networks (GAN)-based ones, represent a huge breakthrough in image generation, due to the ability to generate artificial images from an initial dataset and then make use of them to "predict" image features. GANs are composed of two networks: a *generator* network, that creates tentative fake images, and a *discriminator* network, that aims at identifying whether the generated images are indicative of real-world evidence or not [1]. Nevertheless, GANs are inherently unstable, and suffer from both the lack of meaningful measures to evaluate the quality of their result and limited sample generation capabilities when only a little representative of the population is available.

In the biomedical context, the availability of huge datasets is a major concern: it is indeed a difficult task, as it requires continuous efforts, especially in the long term. Recent advancements in deep learning have shown promising results in classifying pathologies and segmenting medical images [5,9,20,21]. However, these improvements necessitate a consistent and diverse dataset for effective training, highlighting the urgent issue of data scarcity and the need for effective solutions. Image data augmentation techniques address this limitation by generating additional medical images, which can be used to design and refine automated assessment methods for pathological conditions. By doing so, these techniques assist healthcare providers in identifying the most appropriate preventive interventions and therapeutic strategies without relying solely on large medical datasets [12].

Kossen et al. [24] used GANs to create synthetic brain data and corresponding labels, showing good performance in the arterial brain vessel segmentation task. Similarly, Toikkanen et al. [31] used GAN to improve the quality of the predictive model in localizing the hemorrhage from computerized tomography (CT) scans. Synthetic samples from generative models have been demonstrated to alleviate the in-balance and scarcity of labeled training issues. In the same context, Zhai et al. [33] proposed a novel asymmetric semi-supervised GAN (ASSGAN) to

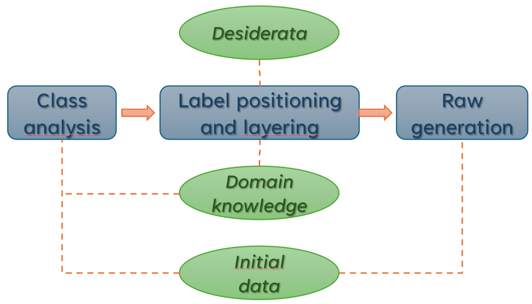

Fig. 1: IDADA Workflow: The process begins with analyzing the classes to identify dynamic and static classes based on the initial data. Next, it computes the label positioning for each dynamic class and generates a comprehensive label encompassing all classes, according to the desired criteria and domain knowledge. Finally, a photo-realistic raw image referencing the initial data is created.

generate reliable segmentation-predicted masks. The authors show that in the absence of labeled data, the network can make use of unlabeled data to improve segmentation performance.

ASP can be used to declaratively express quantitative and qualitative desiderata as well as content generation strategies, providing one with the possibility to easily increment, modify, and update new knowledge at will. Among the different applications, ASP has also been explored in the realm of scheduling, where it can optimize resource allocation and improve operational efficiency [13,14,15]. Moreover, ASP has been applied in the improvement of medical image segmentation and quality (e.g., [6,7]), although these efforts primarily focus on enhancing existing images rather than generating new ones. Other logic-based contributions have emerged in the related field of Content Generation, where, for example, ASP has been employed to produce game content with desirable properties [11,28,30]. To the best of our knowledge, IDADA, a work inspired by [8], represents one of the first logic-based approaches specifically designed for image data augmentation.

3 The IDADA Framework

In this section we introduce IDADA, a framework built upon an Inductive-Deductive Approach for Data Augmentation. We start by illustrating the general workflow; then, we discuss a more detailed application to a specific case study.

3.1 Workflow

The workflow of the proposed framework is illustrated in Figure 1. At the first step, original images are analyzed in order to find all possible featured objects. Thus, a number of classes are identified collecting such objects, and are properly assigned with a label. Note that, depending on the specific scenario, this phase might require the availability of proper domain knowledge. Hence, two kinds of classes are defined, producing a grouping of all possible objects into two categories: *dynamic classes*, featuring objects that can significantly vary from image to image (e.g., size, position, shape, orientation, etc.); and *static classes*, which can be basically considered "fixed" due to their minimal variation across different images.

The second step focuses on dynamic classes, and consists of the generation of new eligible combinations (in terms of position or other spatial features) of objects of such classes. To achieve this, we rely on Answer Set Programming (ASP) to declaratively express how such combinations must be performed. In particular, ad hoc ASP programs produce proper label positioning in images, according to given requirements that can be related to both absolute location and relative positioning/layering of objects of dynamic classes within the image itself. The output is basically a set of (labeled) image structures. Notably, the declarative ASP-based approach allows us to easily integrate the domain knowledge and custom criteria to which all generated image structures and labels must comply.

The last step consists of the generation of new raw images corresponding to the previously created labels (i.e.: if semantically segmented, the raw images match the corresponding labels produced earlier); in particular, the final result is a set of images that: (*i*) differ from the original ones in terms of position of dynamic classes; (*ii*) comply with all the specified criteria (both domain knowledge and custom desiderata); (*iii*) are photo-realistic and can be used for augmenting the original dataset.

3.2 IDADA at work

In the following we illustrate the application of the framework on a case study: the Cataract Dataset for Image Segmentation (CaDIS) [19].

Class Analysis. As introduced above, the first step requires to analyze the dataset with the aim of identifying the classes comprising all elements possibly featured in the images. In general, this would benefit from proper domain knowledge and potentially even interactions with domain experts, for correctly performing semantic segmentation. As for the CaDIS dataset, this information is already available, as each of the 4670 raw frames extracted from 25 videos of cataract surgery is paired with a semantic labelling annotated by experts (from this point onward, they will be referred to as "original labelling" for ease of reading). In particular, images comprise 36 classes: 29 surgical instrument classes, 4 anatomy classes, and 3 miscellaneous classes.

The analysis was completed by determining, for each class, whether it should be considered dynamic; the choice for each class is made based on prior knowledge and input from domain experts. In this case, we checked how classes change across images (in doing this, given that this is a medical domain, we relied both on the labeled images and the raw frames). Specifically, in the CaDIS dataset,

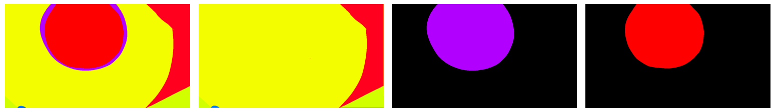

(a) (b) (c) (d)

Fig. 2: Example of class analysis: (a) original label; (b) static classes considered as fixed background; (c) iris and (d) pupil, as dynamic classes, split in single binary segmentation.

the *pupil*, the *iris*, and the *surgical instruments* classes are considered as dynamic classes, while the classes *cornea*, *skin*, *surgical tape*, and *eye retractors* are considered as static ones.

Next, the actual input for the next step in the framework is prepared. In particular, each label corresponding to a dynamic class is extracted to generate a segmentation for that specific class (referred to as binary segmentation). All the binary segmentations of each dynamic class are collected together to create a representative dataset of the possible shapes assumed by that specific object, which will be layered with the other labels afterward to create a new label. Figure 2 shows an example of classification comprising the classes iris and pupil. Specifically, starting from the original label (Figure 2 (a)), the static classes are selected and fixed as background (Figure 2 (b)). In contrast, the dynamic ones, (i.e., iris and pupil) are split into single binary segmentation (Figure 2 (c) and (d), respectively).

Label positioning and layering. At the second step, we make use of an ad-hoc ASP program for managing combinations and layering of labels. The design of this stage can vary significantly, depending on the specific domain and the type of images one aims to generate, to the greatest extent. In this regard, the flexibility of ASP in Knowledge Representation and Reasoning is extremely beneficial.

In the following, we describe how we approach the problem in the chosen domain; we assume that the reader is familiar with the basic notions of the ASP language and refer to [10] for details.

As already mentioned, the main task in this phase consists in the design of proper ASP programs encoding the needed knowledge. For the case of images in the CaDIS dataset, we built an ASP program that takes as input labels of size 540×960 pixels, represented as matrices of the same size. Matrix elements are modeled by facts of the form `cell(R, C, CLASS)`, where variables `R`, `C`, and `CLASS` are mapped to rows and columns of the matrix and the class associated to that cell, respectively. The output produces a new valid positioning of the objects in the image, represented by atoms of the form `new_cell(R, C, CLASS)`, denoting the class `CLASS` that is contained in each cell (`R,C`). In order to effectively express viable choices for the new positions of objects, the idea is to identify a specific number of "pivot points", which can vary depending on the object. These pivot points are then used to accurately reposition the object within the image.

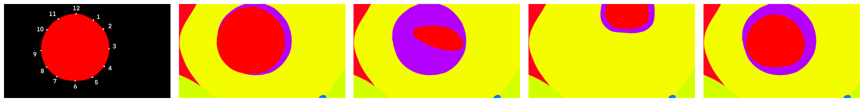

(a) (b) (c) (d) (e)

Fig. 3: Example of positioning and layering of iris and pupil classes: (a) identification of pivot points for the pupil, (b) incorrect layering of the pupil on top of the iris, (c) incorrect shape and proportion of the pupil relative to the iris, (d) incorrect proportion of the iris relative to the cornea, and (e) correct positioning and layering of the pupil and iris.

The repositioning approach takes into account the hierarchical relationships between classes, determining the sequence in which each class should appear based on prior knowledge or desiderata. This also dictates the order in which the final label is reassembled. In the case of discourse, i.e., cataract images, the process begins with the fixed background layer, followed by the addition of the iris layer, and finally the placement of the pupil layer on top. Each label corresponds to a layer that contributes to the final image. For the pupil and iris classes, we opted to identify 12 pivot points situated along the contour of both the pupil and iris classes, thus simplifying the modeling of their shapes and positions. More in detail, we first guess four main pivots, one for each cardinal point, by means of the following choice rules [10]:

```
{pivot(R,C,3,pupil) : max_col(C,pupil), cell(R,C,pupil)} = 1.
{pivot(R,C,6,pupil) : max_row(R,pupil), cell(R,C,pupil)} = 1.
{pivot(R,C,9,pupil) : min_col(C,pupil), cell(R,C,pupil)} = 1.
{pivot(R,C,12,pupil) : min_row(R,pupil), cell(R,C,pupil)} = 1.
```

where predicates `max_col(C,pupil)`, `max_row(R,pupil)`, `min_col(C,pupil)`, and `min_row(R,pupil)` represent the column `C` and the row `R` having the maximum (resp., minimum) index cell containing the class pupil. The remaining eight pivots are then identified as follows: we derive two additional pivots from each pair of adjacent main pivots ($3-6$, $6-9$, $9-12$, $12-3$) by first calculating the difference between the rows and columns of the considered pair; one of the two pivots corresponds to the cell lying on the contour of the pupil with a row equal to the midpoint between the rows of the two main pivots. Similarly, the remaining pivot is calculated using the same reasoning applied to the columns.

We point out that the identification of the pivot points is crucial for understanding, for instance, the object inclination, and thus for selecting (via choice rules) from the representative dataset the appropriate label for the construction of the new image. This process is essential for accurate image generation.

Next, we define the following constraints to ensure that the construction of the image follows the (medical) domain knowledge:

```
(c1) :- pivot(R,C,_,pupil), not cell(R,C,iris).
(c2) :- pivot(R,C,_,iris), not cell(R,C,cornea).
(c3) :- centers_distance(D), D > 70.
(c4) :- width(WP,pupil), width(WI,iris), height(HP,pupil),
        height(HI,iris), WP > HP*WI/HI + 80.
(c5) :- width(WP,pupil), width(WI,iris), height(HP,pupil),
        height(HI,iris), WP < HP*WI/HI - 80.
(c6) :- area(AI,iris), area(AC,cornea), R = AC/AI, R > 4.
```

Here, (c_1) ensures that the pupil lies on the iris. (c_2) forces all of the *iris* class pivots to lie on a cell of the *cornea* class. (c_3) forbids the generation of images where the pupil is not properly centered with respect to the iris. Specifically, the predicate `centers_distance(D)` stands for the distance `D`, expressed in pixel, between the pupil's and the iris' centers. We impose this distance not to exceed an empirically chosen threshold, which amounts to 70 pixel. (c_4) and (c_5) force the iris and the pupil to have similar shapes. The predicates `width` and `height` indicate, for both pupil and iris, the width and height, respectively; for ensuring that their dimensions are proportionate, we impose an empirical tolerance of ± 80 pixels. (c_6) ensures that the proportion between the iris and cornea is correctly observed. We use the predicate `area` to address the area of the iris and cornea; we then force that the cornea area is at most 4 times the iris area. Figure 3 illustrates the application of the approach to the iris and pupil classes. The ASP program generates the twelve pivot points for the pupil class (Figure 3 (a)), ensuring (thanks to the modelled constraints) that the pupil's positioning over the iris aligns with domain knowledge. For instance, Figure 3 (b) depicts a pupil incorrectly positioned on the iris, violating constraint (c_1). Figure 3 (c) shows a pupil with incorrect shape and proportion placed over the iris, violating constraints (c_4) and (c_5). Figure 3 (d) shows an iris with incorrect proportion relative to the cornea, violating constraint (c_6). Finally, Figure 3 (e) demonstrates a correctly generated label.

To manage the several instrument classes with similar properties, we grouped 26 out of the 29 instrument classes into seven categories, and the remaining three classes into class *eye retractors*, which is identified as a fixed class and is already part of the background, whereas the remaining two classes *iris hooks* and *marker* required special pre-processing. Since multiple instances of the same instrument often appear in a single image, we defined a pre-processing used to increase the number of labels depicting the available instruments: in particular, we generated separate binary segmentation for each instance using Python and ASP scripts.

The modeling of surgical instruments follows the strategy described above, including identifying the pivot point(s) and choosing an appropriate binary segmentation to add to the image we are creating.

The interested reader can find all material at `https://github.com/DeMaCS-UNICAL/Data-augmentation-via-ASP`.

Raw generation. At this step, the goal is to create realistic images that comply with the semantically-guided image structure generation. To this aim, in this case we use SPADE (Semantic Image Synthesis With Spatially-Adaptive Normalization) [26], a state-of-the-art paired-data technique, for generating proper textures for each image part. SPADE processes the input semantic layout (defining parts or objects and their spatial relationships) using convolution, normalization, and nonlinearity layers. In SPADE, the mask is projected onto an embedding space and convolved to produce modulation parameters, which are tensors with spatial dimensions. These tensors are multiplied and added to the normalized activation element-wise. We used the same configuration as the original authors [26], i.e., learning rates of 0.0001 for the generator, 0.0004 for the discriminator, and ADAM optimizer, and trained the network for 100 epochs.

4 Experimental analysis

In order to assess the overall effectiveness of the IDADA approach to data augmentation, along with its versatility when dealing with different scenarios, we tested the presented framework over two very diverse real datasets. In particular, besides the already mentioned CaDIS dataset that features cataract images, we considered also the Laryngeal Endoscopic Images dataset [25]. Just like for the CaDIS dataset, also for this case all material is completely available at `https://github.com/DeMaCS-UNICAL/Data-augmentation-via-ASP`.

The Laryngeal Endoscopic Images dataset contains 536 manually segmented *in vivo* color images of the larynx, captured from videos recorded during two resection surgeries. The images include seven classes: *void*, *vocal folds*, *other tissue*, *glottal space*, *pathology*, *surgical tool*, and *intubation*. The dataset includes eight sequences from two patients, categorized into five different groups based on the features exhibited by the images. We divided the classes into two categories based on their relevance. In particular, we designated *vocal folds*, *glottal space*, and *other tissue* as static classes, as their shapes and appearances remain relatively consistent across different images. In contrast, we focused on the *pathology*, *intubation*, and *surgical tool* classes, chosen as dynamic classes.

In both the scenarios of test, IDADA proved to be a viable approach for agumenting the image datasets, granting coherence with background knowledge, compliance with custom desiderata and similarity of the final results. This is particularly encouraging, especially noting that the two case studies significantly differ to a large extent. For a qualitative evaluation of our approach, we showcase selected results in Figure 4. The examples illustrate the effectiveness of our method in generating images that closely match the originals.

For the sake of reproducibility and transparency, all results of our experiments can be found at `https://github.com/DeMaCS-UNICAL/Data-augmentation-via-ASP`.

Furthermore, besides considerations about the viability of the approach and qualitative assessments, we also wanted to perform a more formal analysis. In particular, we wanted to quantitatively evaluate the similarity between the original images and the generated synthetic ones. To this aim, we used the Kernel Inception Distance (KID) [2]. KID is typically used for measuring the quality

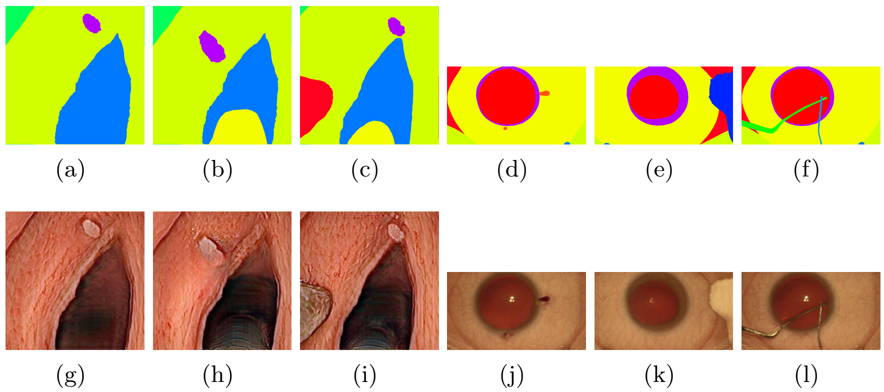

Fig. 4: Examples of laryngeal endoscopic images (a-g,b-h,c-i) and cataract images (d-j,e-k,f-l) obtained by IDADA. Labeled images and the corresponding raw ones are reported in the first and second row, respectively.

of generative models [3] by comparing the distributions of generated and real images using the squared Maximum Mean Discrepancy (MMD) with a polynomial kernel. Lower KID scores indicate higher similarity and better quality of generated images. KID is advantageous for its unbiased nature and reliability, particularly with smaller sample sizes. On the CaDIS dataset, the resulting KID score is 0.10, indicating a very high level of similarity. This proves the quality of the results obtained by means of IDADA, due to both the semantically-guided label generation pervormed via ASP and SPADE's impressive capability to capture the nuanced characteristics of the original data, along with robustness and precision in generating high-fidelity images from labels. Similarly, on the vocal folds dataset, the KID score of 0.26 shows a commendable level of similarity between the generated and original images. While slightly higher than the CaDIS score, this still reflects a strong performance of SPADE, effectively capturing many important features of the original images; on the overall, results underscore its versatility and effectiveness across different datasets. The significant performance on CaDIS, combined with the solid results on vocal folds, suggests that the use of ASP for label generation combined with SPADE for finishing the raw images make IDADA a powerful and reliable tool for generating realistic images from original (potentially limited) datasets.

5 Conclusion and Perspectives

In this paper we presented IDADA, a framework aimed at enabling the declarative specification of data augmentation processes. In particular, we relied on Answer Set Programming (ASP) to guide the generation of realistic images within the biomedical domain. Our methodology involves collecting a dataset of labeled images and identifying the static and dynamic classes to work exclusively with those classes that show a significant change in the images; this allows for more efficient processing and enhances the relevance of the generated images. Using ASP reasoning tasks, we generate new labeled images by detailing the possible appearances of these elements and composing the output appropriately. The newly created semantically labeled images are used as input for proper DL-based methods, which then produce the final pseudo-realistic images.

We assessed the effectiveness of IDADA using images from two distinct datasets, featuring cataract and laryngeal endoscopic images, respectively, yielding promising results. The experiments show that declarative specifications can be seamlessly integrated into the image data augmentation process.

It is worth noting that ASP-encoded specifications allow us to incorporate both domain knowledge and desiderata, providing significant customization capabilities for generating new raw images. The approach eliminates the need for manually finding, collecting, and adapting data within a domain (e.g., surgical images with a specified number of instruments or particular organ positions). Furthermore, ASP grants robustness in handling knowledge updates, for instance allowing easy modifications of logic programs to change the number of elements in a given class or to adjust spatial relationships among elements (e.g., generating images with different numbers of instruments).

On a broader scale, incorporating ASP as the declarative formalism within the data augmentation process enables the collection and translation of declarative specifications (i.e., logic programs) into properly generated labeled images suitable for DL methods. With this respect, SPADE, in particular, has shown satisfactory results, producing realistic raw images that accurately correspond to their labeled counterparts.

Future work will focus on conducting experimental campaigns to evaluate the quality of the generated images in relation to specific tasks in the biomedical domain, and to assess the performance of our approach on additional datasets.

Acknowledgements

This work has been partially supported by: PNRR project FAIR - Future AI Research (PE00000013), Spoke 9 - Green-aware AI, under the NRRP MUR program funded by the "NextGenerationEU"; PNRR project Tech4You "Technologies for climate change adaptation and quality of life improvement", CUP H23C22000370006, under the NRRP MUR program funded by the "NextGenerationEU"; PON “Ricerca e Innovazione” 2014-2020, CUP: H25F21001230004. Francesco Calimeri is member of the Gruppo Nazionale Calcolo Scientifico-Istituto Nazionale di Alta Matematica (GNCS-INdAM).

References

1. Aggarwal, A., Mittal, M., Battineni, G.: Generative adversarial network: An overview of theory and applications. International Journal of Information Management Data Insights **1**(1), 100004 (2021)

2. Bińkowski, M., Sutherland, D.J., Arbel, M., Gretton, A.: Demystifying mmd gans. arXiv preprint arXiv:1801.01401 (2018)
3. Böhland, M., Bruch, R., Löffler, K., Reischl, M.: Unsupervised gan epoch selection for biomedical data synthesis. In: Current Directions in Biomedical Engineering. vol. 9, pp. 467–470. De Gruyter (2023)
4. Brewka, G., Eiter, T., Truszczynski, M.: Answer set programming at a glance. Commun. ACM **54**(12), 92–103 (2011). `https://doi.org/10.1145/2043174.2043195`, `http://doi.acm.org/10.1145/2043174.2043195`
5. Bruno, P., Calimeri, F., Kitanidis, A.S., De Momi, E.: Understanding automatic diagnosis and classification processes with data visualization. In: 2020 IEEE international conference on human-machine systems (ICHMS). pp. 1–6. IEEE (2020)
6. Bruno, P., Calimeri, F., Marte, C.: Dedudeep: An extensible framework for combining deep learning and asp-based models. In: Logic Programming and Nonmonotonic Reasoning: 16th International Conference, LPNMR 2022, Genova, Italy, September 5–9, 2022, Proceedings. pp. 505–510. Springer (2022)
7. Bruno, P., Calimeri, F., Marte, C., Manna, M.: Combining deep learning and asp-based models for the semantic segmentation of medical images. In: Rules and Reasoning: 5th International Joint Conference, RuleML+ RR 2021, Leuven, Belgium, September 13–15, 2021, Proceedings 5. pp. 95–110. Springer (2021)
8. Bruno, P., Calimeri, F., Marte, C., Perri, S.: Data augmentation: a combined inductive-deductive approach featuring answer set programming. arXiv preprint arXiv:2310.14413 (2023)
9. Bruno, P., Spadea, M.F., Scaramuzzino, S., De Rosa, S., Indolfi, C., Gargiulo, G., Giugliano, G., Esposito, G., Calimeri, F., Zaffino, P.: Assessing vascular complexity of paod patients by deep learning-based segmentation and fractal dimension. Neural Computing and Applications **34**(24), 22015–22022 (2022)
10. Calimeri, F., Faber, W., Gebser, M., Ianni, G., Kaminski, R., Krennwallner, T., Leone, N., Maratea, M., Ricca, F., Schaub, T.: Asp-core-2 input language format. Theory Pract. Log. Program. **20**(2), 294–309 (2020). `https://doi.org/10.1017/S1471068419000450`, `https://doi.org/10.1017/S1471068419000450`
11. Calimeri, F., Germano, S., Ianni, G., Pacenza, F., Pezzimenti, A., Tucci, A.: Answer set programming for declarative content specification: A scalable partitioning-based approach. In: Ghidini, C., Magnini, B., Passerini, A., Traverso, P. (eds.) AI*IA 2018 - Advances in Artificial Intelligence - XVIIth International Conference of the Italian Association for Artificial Intelligence, Trento, Italy, November 20-23, 2018, Proceedings. Lecture Notes in Computer Science, vol. 11298, pp. 225–237. Springer (2018). `https://doi.org/10.1007/978-3-030-03840-3_17`, `https://doi.org/10.1007/978-3-030-03840-3_17`
12. Chen, Y., Yang, X.H., Wei, Z., Heidari, A.A., Zheng, N., Li, Z., Chen, H., Hu, H., Zhou, Q., Guan, Q.: Generative adversarial networks in medical image augmentation: a review. Computers in Biology and Medicine p. 105382 (2022)
13. Dodaro, C., Galatà, G., Gebser, M., Maratea, M., Marte, C., Mochi, M., Scanu, M.: Operating room scheduling via answer set programming: improved encoding and test on real data. Journal of Logic and Computation p. exae041 (2024)
14. Dodaro, C., Galatà, G., Marte, C., Maratea, M., Mochi, M.: Nuclear medicine scheduling via answer set programming (2024)
15. Dodaro, C., Maratea, M.: Nurse scheduling via answer set programming. In: Logic Programming and Nonmonotonic Reasoning: 14th International Conference, LPNMR 2017, Espoo, Finland, July 3-6, 2017, Proceedings 14. pp. 301–307. Springer (2017)

16. Durgadevi, M., et al.: Generative adversarial network (gan): a general review on different variants of gan and applications. In: 2021 6th International Conference on Communication and Electronics Systems (ICCES). pp. 1–8. IEEE (2021)
17. Eiter, T., Ianni, G., Krennwallner, T.: Answer set programming: A primer. In: Tessaris, S., Franconi, E., Eiter, T., Gutiérrez, C., Handschuh, S., Rousset, M., Schmidt, R. (eds.) Reasoning Web. Semantic Technologies for Information Systems, 5th International Summer School 2009, Brixen-Bressanone, Italy, August 30 - September 4, 2009, Tutorial Lectures. Lecture Notes in Computer Science, vol. 5689, pp. 40–110. Springer (2009). https://doi.org/10.1007/978-3-642-03754-2_2, https://doi.org/10.1007/978-3-642-03754-2_2
18. Gelfond, M., Lifschitz, V.: Classical negation in logic programs and disjunctive databases. New Generation Computing **9**(3/4), 365–386 (1991). https://doi.org/10.1007/BF03037169, https://doi.org/10.1007/BF03037169
19. Grammatikopoulou, M., Flouty, E., Kadkhodamohammadi, A., Quellec, G., Chow, A., Nehme, J., Luengo, I., Stoyanov, D.: Cadis: Cataract dataset for surgical rgb-image segmentation. Medical Image Analysis **71**, 102053 (2021)
20. Huang, Y., Yang, X., Liu, L., Zhou, H., Chang, A., Zhou, X., Chen, R., Yu, J., Chen, J., Chen, C., et al.: Segment anything model for medical images? Medical Image Analysis **92**, 103061 (2024)
21. Jiang, H., Diao, Z., Shi, T., Zhou, Y., Wang, F., Hu, W., Zhu, X., Luo, S., Tong, G., Yao, Y.D.: A review of deep learning-based multiple-lesion recognition from medical images: classification, detection and segmentation. Computers in Biology and Medicine **157**, 106726 (2023)
22. Kaufmann, B., Leone, N., Perri, S., Schaub, T.: Grounding and solving in answer set programming. AI Magazine **37**(3), 25–32 (2016)
23. Khalifa, N.E., Loey, M., Mirjalili, S.: A comprehensive survey of recent trends in deep learning for digital images augmentation. Artificial Intelligence Review **55**(3), 2351–2377 (2022)
24. Kossen, T., Subramaniam, P., Madai, V.I., Hennemuth, A., Hildebrand, K., Hilbert, A., Sobesky, J., Livne, M., Galinovic, I., Khalil, A.A., et al.: Synthesizing anonymized and labeled tof-mra patches for brain vessel segmentation using generative adversarial networks. Computers in biology and medicine **131**, 104254 (2021)
25. Laves, M.H., Bicker, J., Kahrs, L.A., Ortmaier, T.: A dataset of laryngeal endoscopic images with comparative study on convolution neural network-based semantic segmentation. International journal of computer assisted radiology and surgery **14**(3), 483–492 (2019)
26. Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2337–2346 (2019)
27. Roß, T., Bruno, P., Reinke, A., Wiesenfarth, M., Koeppel, L., Full, P.M., Pekdemir, B., Godau, P., Trofimova, D., Isensee, F., et al.: Beyond rankings: Learning (more) from algorithm validation. Medical image analysis **86**, 102765 (2023)
28. Shaker, N., Togelius, J., Nelson, M.J.: Procedural Content Generation in Games. Computational Synthesis and Creative Systems, Springer (2016). https://doi.org/10.1007/978-3-319-42716-4, https://doi.org/10.1007/978-3-319-42716-4
29. Shorten, C., Khoshgoftaar, T.M.: A survey on image data augmentation for deep learning. Journal of big data **6**(1), 1–48 (2019)
30. Smith, A.M., Mateas, M.: Answer set programming for procedural content generation: A design space approach. IEEE Trans. Comput. Intell. AI Games **3**(3),

187–200 (2011). https://doi.org/10.1109/TCIAIG.2011.2158545, https://doi.org/10.1109/TCIAIG.2011.2158545
31. Toikkanen, M., Kwon, D., Lee, M.: Resgan: Intracranial hemorrhage segmentation with residuals of synthetic brain ct scans. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part I 24. pp. 400–409. Springer (2021)
32. Wu, H., Liu, Q., Liu, X.: A review on deep learning approaches to image classification and object segmentation. Computers, Materials & Continua **60**(2) (2019)
33. Zhai, D., Hu, B., Gong, X., Zou, H., Luo, J.: Ass-gan: Asymmetric semi-supervised gan for breast ultrasound image segmentation. Neurocomputing **493**, 204–216 (2022)