

Article

Cancer Type Classification in Liquid Biopsies Based on Sparse Mutational Profiles Enabled through Data Augmentation and Integration

Alexandra Danyi ¹, Myrthe Jager ^{1,2} and Jeroen de Ridder ^{1,2,*}

¹ Center for Molecular Medicine, University Medical Center Utrecht, 3584 CX Utrecht, The Netherlands; A.Danyi@umcutrecht.nl (A.D.); M.Jager-2@umcutrecht.nl (M.J.)

² Oncode Institute, 3521 AL Utrecht, The Netherlands

* Correspondence: J.deRidder-4@umcutrecht.nl

Abstract: Identifying the cell of origin of cancer is important to guide treatment decisions. Machine learning approaches have been proposed to classify the cell of origin based on somatic mutation profiles from solid biopsies. However, solid biopsies can cause complications and certain tumors are not accessible. Liquid biopsies are promising alternatives but their somatic mutation profile is sparse and current machine learning models fail to perform in this setting. We propose an improved method to deal with sparsity in liquid biopsy data. Firstly, data augmentation is performed on sparse data to enhance model robustness. Secondly, we employ data integration to merge information from: (i) SNV density; (ii) SNVs in driver genes and (iii) trinucleotide motifs. Our adapted method achieves an average accuracy of 0.88 and 0.65 on data where only 70% and 2% of SNVs are retained, compared to 0.83 and 0.41 with the original model, respectively. The method and results presented here open the way for application of machine learning in the detection of the cell of origin of cancer from liquid biopsy data.



Citation: Danyi, A.; Jager, M.; de Ridder, J. Cancer Type Classification in Liquid Biopsies Based on Sparse Mutational Profiles Enabled through Data Augmentation and Integration. *Life* **2022**, *12*, 1. <https://doi.org/10.3390/life12010001>

Academic Editors: Lefteris Koumakis, George Potamias and Alexandros Kanterakis

Received: 27 November 2021

Accepted: 17 December 2021

Published: 21 December 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Keywords: deep learning; genomics; genetic variability; bioinformatics

1. Introduction

Identification of the ‘cell of origin’ of cancer—which is representative of its anatomical origin and histological characteristics—is an important step in effective cancer treatment, as it has been established that the cell of origin can be a strong predictor of response to therapy and overall prognosis [1,2]. Importantly, tailoring targeted treatments to the cell of origin can result in more successful treatment and better prognosis for ‘cancer of unknown primary’ (CUP) patients, where the cell of origin is often difficult to determine with standard histopathology techniques [3]. One particularly promising alternative to identify the cell of origin in CUP patients is based on somatic mutations derived from whole genome sequencing (WGS) data. Recent studies demonstrate proof of concept of this, by leveraging machine learning to accurately identify the cell of origin using somatic mutation density profiles obtained in solid biopsies of primary and metastatic cancers [4,5]. However, in many cases taking solid tissue biopsies is challenging, as accessibility to the tumor can be limited. Moreover, conventional needle biopsies are invasive and may lead to clinical complications in one out of six procedures [6]. Even in case a biopsy is feasible, solid tissue biopsies cannot be used for early detection and screening since the location of the tumor is not yet known by definition.

Blood plasma of cancer patients contains circulating tumor DNA (ctDNA) which can be used to detect somatic mutations [7–9]. Liquid biopsies of the blood may, therefore, provide a valuable alternative to solid biopsies. Moreover, conventional solid tissue biopsy can only capture tumor heterogeneity to a limited extent as only one spatial location of the tumor is sampled, while ctDNA originates from the whole tumor and thus captures more of the tumor heterogeneity [10,11]. Another advantage of liquid biopsies is that

on top of common solid biopsy biomarker modalities (such as somatic mutations), the fragmentation pattern of cell-free DNA (cfDNA) can also be leveraged for diagnostic purposes [12–14]. Coupling fragmentation and mutational profile data to classify the cell of origin of the tumor is a promising diagnostic avenue. In [14], a machine learning approach was proposed to detect cancer using genome-wide fragmentation data combined with mutation analysis in early stage cancer liquid biopsy samples. Furthermore, the benefits of exploiting multiple liquid biopsy analytes (LBA) were previously shown, where using data from matched circulating tumor cell (CTC) mRNA, CTC genomic DNA (gDNA), extracellular vesicle (VE) mRNA and cfDNA enhanced the detection of clinically actionable information [15]. A major challenge for ctDNA-based cancer diagnostics, however, is that ctDNA only represents a small fraction of the total cfDNA. As a result, the amount of tumor DNA that can be captured in liquid biopsies is much lower than in conventional solid tissue biopsies. While in patients with advanced-stage cancers the ctDNA concentration can exceed 10% of the total cfDNA (thus increasing the chance of detecting comprehensive somatic mutation profiles), it is significantly lower in patients with cancer types such as glioma, medulloblastoma, bladder, renal and gastroesophageal cancer [10]. Moreover, ctDNA levels are likely to be lower in early-stage tumors where the tumor burden is low [7,16]. While ctDNA levels can vary depending on tumor type and disease stage, the generally limited amount of ctDNA in liquid biopsies results in a much sparser somatic mutation profile when compared to solid tissue biopsies.

Liquid biopsies are only a viable alternative to solid biopsies if there is a certain level of concordance between the somatic mutations obtained from ctDNA and the tissue, which is generally high in case of metastatic tumors. According to a prospective study on advanced metastatic cancers, 72% of somatic SNV/indel mutations found in solid tissue biopsy samples were detected by the applied liquid biopsy assay with high-intensity sequencing [17] and similarly in an earlier study on metastatic prostate and breast cancers, the average concordance was determined as 88% [18]. However, in a recent study on stage I-IV non-small lung cell cancer (NSCLC) patients, the percentage of concordant mutations between cfDNA and tissue biopsy ranged between only 12.4% (stage I) and 73.8% (stage IV) [19]. This shows that lower ctDNA concentrations in early stage cancer can result in a more dramatic sparsity which might complicate the utilization of somatic mutations to identify the cell of origin, while metastatic cancers with a higher tumor load are less affected by this issue. Therefore, there is no generally defined concordance level of somatic mutations in a liquid biopsy when compared to a corresponding solid biopsy sample, but it is clear that the variable concordance of somatic mutations poses a challenge to diagnostic applications due to potentially high levels of sparsity.

Despite the sparsity of somatic mutation profiles in liquid biopsies, efforts have already been made to utilize liquid biopsy data for the identification of the cell of origin of cancer. The majority of these studies successfully apply targeted approaches [7,20], which mainly focus on mutation density in cancer driver genes as input features for cancer type classification. Due to the low abundance of ctDNA, many driver mutations can be missed due to sampling artifacts in WGS, therefore targeted sequencing is necessary to ensure detection sensitivity. However, adding genome-wide information could improve the classification. According to a recent study on genome-wide liquid biopsies of postoperative early stage residual cancers, the integration of genome-wide mutation data allowed sensitive residual disease detection by overcoming the limitations of sparsity [21]. Additionally, previous work on WGS data from solid tissue biopsies have already established that somatic SNV density at 1 Mb scale is the most prominent predictor of cancer type as it represents the genomic imprint of the cell of origin chromatin organization, with passenger somatic SNVs being the most prominent contributors [4,5]. However, the performance of a support vector machine classifier based on passenger SNVs was considerably reduced on artificially generated sparse data [4]. This indicates that current state-of-the-art classifiers might not be robust enough to reliably predict tumor type from somatic mutation profiles which are extremely sparse, such as those obtained from liquid biopsies.

In our work, we explore the utilization of sparse genome-wide somatic mutation data in the classification of the cell of origin of cancer. High quality genome-wide somatic mutation data obtained from ctDNA is very scarce [21], and by no means sufficient to support the training of robust classifiers. Therefore sparse SNV samples are generated based on WGS of primary cancer samples from the PCAWG dataset [22] to model ctDNA conditions. Then, in order to address the challenge of sparsity and enable model robustness, we propose to improve on a state-of-the-art classification approach in two ways. First of all, we recognize that, in a sparse somatic data setting, massive training set sizes are required to enable the classifier to become robust against feature sparsity (i.e., missing somatic mutations), therefore we apply extensive data augmentation. Secondly, to complement the genome-wide somatic SNV density features, we propose to leverage data integration of two additional features derived from the SNV data: trinucleotide counts and mutation density on driver genes. Mutational signatures, which differ between cancer types and reflect the result of diverse oncogenic processes, are constructed from the whole mutation spectrum that consists of trinucleotide (e.g., ACA > AGA) mutation frequencies [23]. These trinucleotide features were utilized for cancer type classification earlier [5]. Trinucleotide mutational contexts are composite features and therefore can be less affected by scarcely available SNV mutations. Additionally, while frequently mutated driver genes can represent one of the major hallmarks of cancer and are commonly shared between cancer types [24], certain combinations of affected driver genes can differ between different tissues [25], which has been used as a basis for cancer type classification before [7] and driver gene mutation density features were found to have some predictive power in a few cancer types (for instance Panc-AdenoCA) [5].

Our results demonstrate that both data augmentation as well as data integration considerably increase the robustness of the classifier. As a result, when the original solid tissue biopsy samples are downsampled to retain only 70% of somatic mutations (an approximation of general sparsity in liquid biopsy of late stage and metastatic cancers [17,19]), the deep learning model proposed here reaches classification accuracies similar to previous state-of-the-art classifiers on solid biopsy (i.e., non-sparse) WGS mutation profiles. These findings raise the possibility that somatic mutation profiles derived from liquid biopsies can aid the detection of the cell of origin in CUP patients.

2. Materials and Methods

2.1. Sparse Data Generation

Somatic SNV consensus calls of 2374 primary tumors based on WGS with an average coverage of ~30X were downloaded in VCF format from the ICGC Data Portal (<http://dcc.icgc.org/releases/PCAWG/>, accessed on 16 January 2018). The 2374 primary tumor sample set was constructed from cancer types where at least 30 samples were available (Supplementary Table S1). The list of used PCAWG sample and donor identifiers can be found in Supplementary Table S2.

In order to model the sparse somatic mutation profile of ctDNA in liquid biopsy, we randomly subsampled 2%, 5%, 10%, 25% and 70% of the somatic SNVs from each primary cancer VCF file. The resulting random samples were then used to generate genome-wide SNV bin density, genome-wide trinucleotide and driver gene SNV density data for each sparse sample. A schematic overview of sparse sample generation can be seen in Figure 1a.

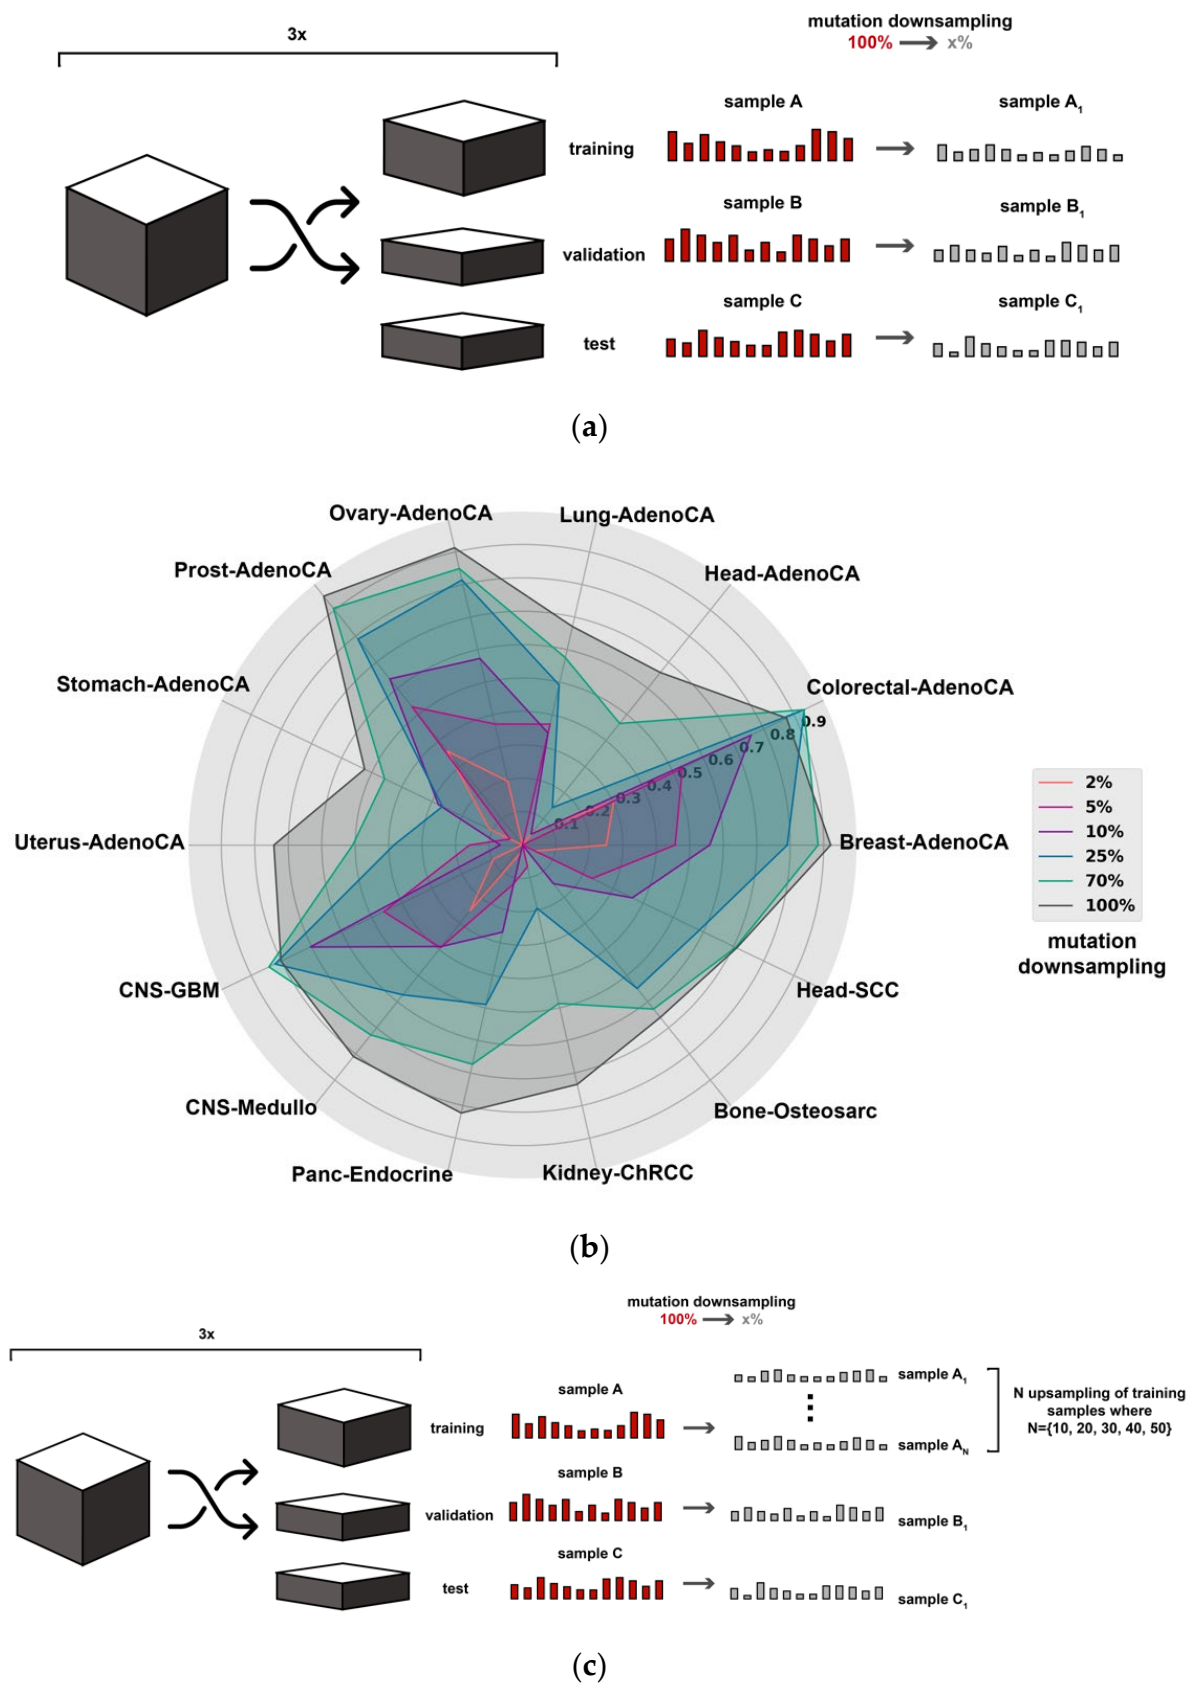


Figure 1. Cont.

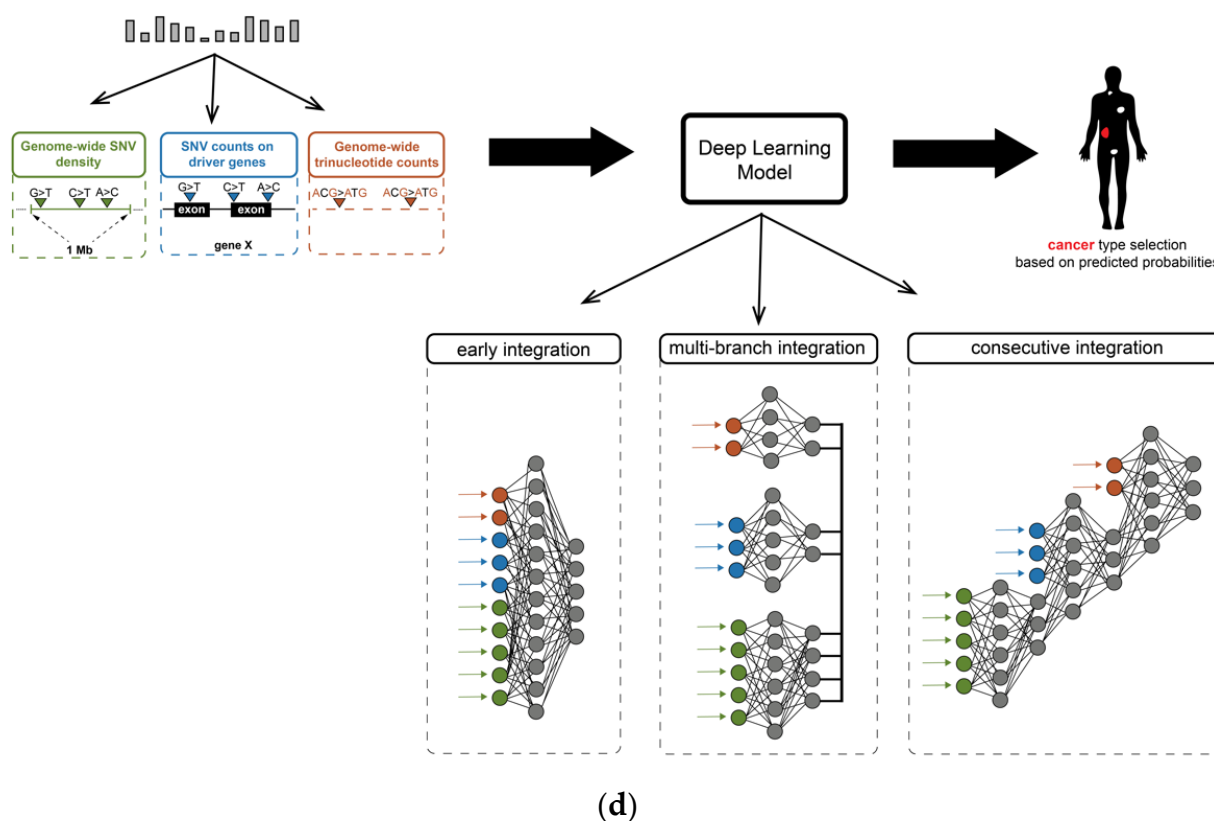


Figure 1. State-of-the-art model performance on the generated sparse datasets and overview of proposed model improvements. (a) Schematic overview of sparse somatic SNV data generation. Sparse samples were generated by retaining 70%, 25%, 10%, 5% and 2% of somatic mutations from each original sample (mutation downsampling) for the training, validation and test sets. Results were evaluated in 3-fold stratified shuffle-split cross-validation. (b) Cancer types classification accuracies of somatic SNV data with different (70%, 25%, 10%, 5%, 2%) mutation downsampling rates and full somatic SNV data (100%). The axes represent the accuracy for the cancer types, expressed in average F1 scores (stratified shuffle split cross-validation with three partitions). Only cancer types with the lowest sparse data accuracies are shown. (c) Schematic overview of data augmentation. Sparse somatic SNV samples (subsamples of each original WGS sample) were generated by retaining 70%, 25%, 10%, 5% and 2% of SNVs from each original sample (mutation downsampling). Data augmentation was applied solely on the training set, where instead of one subsample, $N \in (10, 20, 30, 40, 50)$ different subsamples were generated based on each full somatic SNV sample (sample upsampling). Results were evaluated in 3-fold shuffle-split cross-validation. (d) Schematic overview of data integration of trinucleotide motifs and driver gene mutation densities to tackle sparse somatic SNV data. The integration of the three feature sets was tested in different deep learning architectures: early integration, multi-branch integration and consecutive integration.

2.2. Baseline/Original Model and Additional Implementation

The original classifier [5] was used from <https://github.com/ICGC-TCGA-PanCancer/TumorType-WGS/tree/master/DNN-Model> (accessed on 1 March 2020) and retrained on the available PCAWG samples. In all consecutive steps, this model was used as a basis and further adjustments were applied on it in the data integration experiments. All additional models were implemented and trained in Tensorflow 1.10.0 and Keras 2.1.5. All code was written in Python 3.6.

We used shuffle split cross-validation with three splits, by randomly selecting 60% of the samples for training and 20–20% for validation and testing in each split round, respectively. Stratification was applied in each split, in order to distribute the cancer types according to the relative percentages in the training (60%), validation (20%) and test

sets (20%). Hyperparameters were set by optimizing the performance on the validation set in each shuffle-split partition. The optimized hyperparameter set was the following: learning rate for Adam [26], L2-regularisation penalty (otherwise known as weight decay), dropout rate, the number of hidden layers, the number of neurons per hidden layer and activation function [5]. We used the same Bayesian method and library for hyperparameter optimization as in [5], where the ‘gp_minimize’ function from the scikit-optimize 0.7.4 library was applied [27]. Briefly, the models in each fold were trained using the Adam optimizer with a batch size of 32 for 50 epochs. Bias values were initialised as 0 and all other network weights were initialised using a glorot uniform distribution [28]. Each model was evaluated with 200 hyperparameter combinations that were obtained from the Bayesian optimization process (i.e., ‘gp_minimize’ function was called consecutively 200 times).

In the multi-branch and consecutive data integration experiments, the following hyperparameters were optimized separately for each feature type: weight decay, number of dense layers, number of dense nodes per layer, dropout and activation function. This meant that for the multi-branch experiments, hyperparameters were optimized per branch, while in the consecutive data integration experiments the hyperparameters were set separately after each input entry point. The final hyperparameter settings of each model are described in Supplementary Table S3.

2.3. Accuracy Calculation

Cancer type classification accuracies were calculated as F1-scores to take class imbalance into account, and averaged across the three shuffle split runs. This cross-validation setup was applied in all further experiments when comparing classification accuracy at cancer type level. Additionally, average model accuracy was calculated on the basis of correctly classified samples in all cancer types. F1-scores and confusion matrices were calculated with scikit-learn 0.23.2 [29].

2.4. Data Augmentation

N subsamples were generated from each sample, where $N \in (10, 20, 30, 40, 50)$, N corresponds to the data augmentation level of the dataset which we refer to as ‘10×’, ‘20×’, ‘30×’, ‘40×’ and ‘50×’ (Supplementary Table S4). Sampled mutations were randomly selected for all N subsamples as described in ‘Sparse data generation’. If a subsampled sample was present in the training set, none of its other subsamples was part of the test/validation sets during the classification process.

2.5. Data Integration

The 96 trinucleotide motif features were generated on a per motif basis where all occurrences of a motif (e.g., ACA > AGA) were aggregated, resulting in a summed value for the whole genome in each patient sample.

For the construction of driver gene features, we used the driver gene list released in one of the flagship papers of the latest the Pan-Cancer Atlas release [30]. SNVs were counted on gene regions from the start position of the first exon until the end position of the last exon.

2.6. Feature Importance Assessment

The original published code (<https://github.com/ankurtaly/Integrated-Gradients>, accessed on 1 August 2020) from [31] was used to calculate the integrated gradients for the early integration models. Feature attributions were calculated in respect to the top predicted label in each sample. When calculating integrated gradients, first the gradient of the prediction output is calculated with respect to the feature values of the input in a range of interpolations between the original input (a vector containing all input features) and a baseline. In our case, zero baseline was used for all features. We calculated the feature attributions for each test split.

As this approach can be applied to explain feature attributions on a single sample level, the feature attributions were averaged from all test samples in each cancer type in order to get cancer type specific feature attributions. In order to aid data visualization when comparing the full sparse feature matrices, log-modulus transformation $L(f_i) = \text{sign}(f_i) * \log(|f_i| + 1)$ was applied on the raw feature importance values, where f_i is the feature attribution value of feature f in cancer type i . Then, we averaged the log-modulus transformed attributions of all test samples in each cancer type in order to get the transformed attributions per cancer type. Finally, scaling was applied to each feature matrix individually: $f_{scaled} = \frac{f_i}{f_{max}}$ if $f_i > 0$ and $f_{scaled} = \frac{f_i}{f_{min}}$ if $0 > f_i$. For visualization purposes, feature attributions were sorted per cancer type in the full solid biopsy SNV data and then the sparse data matrices were sorted based on the obtained feature order in the full solid biopsy SNV data, in order to follow the same feature order per cancer type in all matrices of a given feature type. Features that consistently had 0 attribution across all cancer types were excluded from visualization and are not displayed in the heatmaps.

3. Results

3.1. Assessment of State-of-the-Art Deep Learning Model on Sparse Somatic Mutation Data

We first assessed the classification performance of a state-of-the-art deep learning model proposed in [5] in a liquid biopsy setting (Section 2). This model is based on genome-wide SNV density features across 2897 bins where each bin represents a 1 Mb region of the genome. Sparse mutation profiles were artificially generated based on the PCAWG data collection [22], a dataset that consists of 2374 primary cancer samples across 16 cancer types with Soft-Messarc having the smallest sample size of 33 samples and Liver-HCC representing the biggest cancer type with 304 samples (see Supplementary Table S1 for tumor type abbreviations). To model the somatic mutation sparsity of liquid biopsy samples, sparse samples were obtained by selecting a subset of all somatic SNVs (mutation downsampling) from each PCAWG sample. Multiple sparse datasets were generated by retaining 70%, 25%, 10%, 5% and 2% of all somatic SNVs (Figure 1a). The first measurement was designated at 70%, as this value is an approximation of the reported percentage of somatic mutations that can be detected in liquid biopsies compared to solid tissue biopsies in most advanced and metastatic cancers [17,19]. Lower downsampling levels were chosen to mimic lower ctDNA concentrations, which have been observed in certain cancer types and in case of low tumor burden [7,10]. The full somatic SNV data of all 2374 samples was also classified in order to have baseline accuracies of the original classification setting from [5] for each cancer type.

As an initial test, predictions were made on the sparse datasets with the original model trained on the full (100%) somatic SNV data. As expected, accuracy dropped dramatically in case of each sparse dataset, except at 70% (Supplementary Figure S1). When the model is trained and tested on the full 100% SNV data, our results are in line with previous findings [5] and show that for all but five cancer types the classification accuracy exceeds 0.7, and for most cancer types is well above 0.8 (Supplementary Figures S1 and S2, Figure 1b). Stomach-AdenoCA is the worst performing tumor type and is often misclassified as Panc-AdenoCA or Eso-AdenoCA due to similar biological origin, as described in the earlier study [5].

In the next step, we aimed to assess the effect of SNV downsampling on the training procedure and retrained the original model on each sparse dataset, instead of applying the model trained on the full data throughout all sparse test datasets. Overall, when the model is trained and tested on sparse data, a better accuracy trend is appreciable compared to the results when the model is trained on the full somatic SNV data but tested on sparse data (Supplementary Figures S1 and S2, Figure 1b). Generally, when the downsampling rate increases, classification accuracy decreases, often substantially. When only 2% of somatic SNVs are retained, an average accuracy below 0.2 in 10 out of 16 tumor types was observed (Supplementary Figure S2). Still, for six of the cancer types a classification accuracy exceeding 0.5 was retained, suggesting that even at this sparsity level some signal

is retained that can be used by the classifier. Interestingly, for ColoRect-AdenoCA the classification accuracy improved in the 70% downsampled mutation experiment compared to the 100% baseline (Figure 1b). This might suggest that downsampling the mutations has a regularization effect on the model for this particular cancer type. Similarly, the average prediction accuracy of CNS-nonDiffGBM and Lung-SCC also slightly improved at 70% (Supplementary Figure S2).

In summary, the classification accuracy gradually drops as the number of retained somatic mutations decreases, with the least but still considerable performance decrease at 70%, even if the classification model was trained on the corresponding sparse data set (Figure 1b, Supplementary Figure S2). This indicates that the current model is not suitable for liquid biopsy data that has a sparser somatic profile than conventional solid tissue biopsies.

3.2. Data Augmentation

Data augmentation is a technique that was introduced in the field of deep learning to aid image classification problems where the available training data is limited. Essentially, data augmentation enhances the available data in terms of size, for instance through basic sample (image) manipulations (geometrical or color transformations, random erasing, image mixing etc.) or model-based augmentation methods (GANs, adversarial training etc.) [32]. Data augmentation addresses the issue of overfitting during the training process, as the augmented data represents a much more comprehensive collection of possible data points. Therefore, better model generalization and robustness can be achieved. By applying different data augmentation approaches, the error rate of the model can be reduced considerably, even up to 12.1% [33–36].

Contrary to traditional machine learning techniques, the performance of deep learning models keeps improving as the training data size is increasing, according to a power-law relationship [37]. Hence, we hypothesized that reduced classification performance due to sparsity of somatic mutation profiles obtained from liquid biopsies can be mitigated by substantially increasing the training set size. By providing the classifier with more and diverse training data it should have more information to learn robust feature patterns under sparse data conditions. However, since genome-wide somatic mutation data from liquid biopsy of more than a few samples per cancer type is not yet readily available, we investigated if the concept of ‘in silico’ data augmentation can be applied. To achieve sufficiently large training set sizes, augmented training samples were generated based on multiple random downsampling of somatic SNVs in PCAWG WGS samples (Figure 1c).

Multiple data augmentation experiments (generating N subsamples from each original sample, where $N \in (10, 20, 30, 40, 50)$; see section ‘Data augmentation’ in Section 2) were performed at each mutation downsampling rate in order to identify a potential inflexion point after which further data augmentation does not improve the performance (Figure 1c). The different data augmentation rates were compared based on average accuracy (average ratio of correctly classified samples in stratified shuffle split cross-validation with three partitions). Additionally, paired t -tests were performed to determine whether the average accuracy changes were significant (Supplementary Tables S5 and S6, Supplementary Figure S5a). Starred comparisons indicate significant differences at $p < 0.05$. For all downsampling rates, $10\times$ data augmentation greatly improves the performance of the model to an average accuracy of 0.51** compared to 2% baseline with 0.41, 0.66** compared to the 5% baseline with 0.56, 0.73** compared to the 10% baseline with 0.63, 0.81* compared to 25% baseline with 0.75 and 0.87* compared to the 70% baseline with 0.83, respectively (Figure 2a). At the lowest mutation downsampling rates (when 25% and 70% of all SNVs were retained), performance tops with using only $20\times$ data augmentation, at an average accuracy of 0.84* compared to 25% baseline with 0.75 and 0.88* compared to the 70% baseline with 0.83, respectively. At the highest mutation downsampling rates (when retaining only 2% and 5% of all SNVs), the application of $50\times$ data augmentation shows the most considerable average performance improvement compared to the baseline experiments

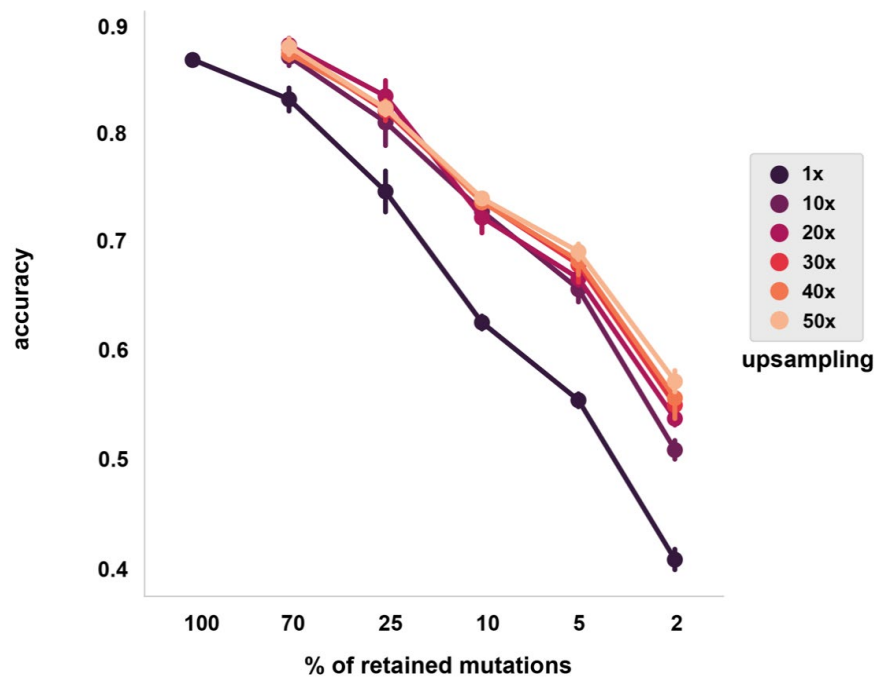
using the original model without data augmentation, at an average accuracy of 0.57** compared to 2% baseline with 0.41 and 0.69*** compared to the 5% baseline with 0.56, respectively (Figure 2a).

The model performance was also assessed at cancer type level in all mutation downsampling experiments (expressed in average F1 scores) (Figure 2b, Supplementary Figure S3). When only 70% and 25% of somatic mutations are retained, data augmentation improves the classifier performance for all cancer types, except ColoRect-AdenoCA, while at 10% mutation downsampling only the accuracy of Head-SCC does not improve. At higher downsampling rates (retaining only 2% and 5% of SNVs) data augmentation improves the classifier performance for all cancer types, except for Uterus-AdenoCA, Lymph-BNHL and Lung-SCC at 5% and Head-AdenoCA and Lung-SCC at 2%. While Lymph-BNHL and Lung-SCC are both classified with high accuracy across all downsampling rates even without data augmentation, Uterus-AdenoCA and Head-AdenoCA show a rather dramatic drop in classification accuracy as the level of sparsity increases and contrary to other cancer types, data augmentation remained ineffective at 2% and 5% downsampling for these two types (Supplementary Figure S2, Figure 2b). Lower classification accuracy can also be partially associated with lower sample numbers in a few cancer types (Figure 2b). When retaining only 2% and 5% of somatic SNVs, all samples are persistently misclassified from a few cancer types without data augmentation (Figure 2b). These results clearly show that data augmentation is fundamental when dealing with classification of the cell of origin of cancer from sparse genome-wide mutation data and it can considerably improve the classification accuracy even at higher downsampling rates (2%, 5%, 10%).

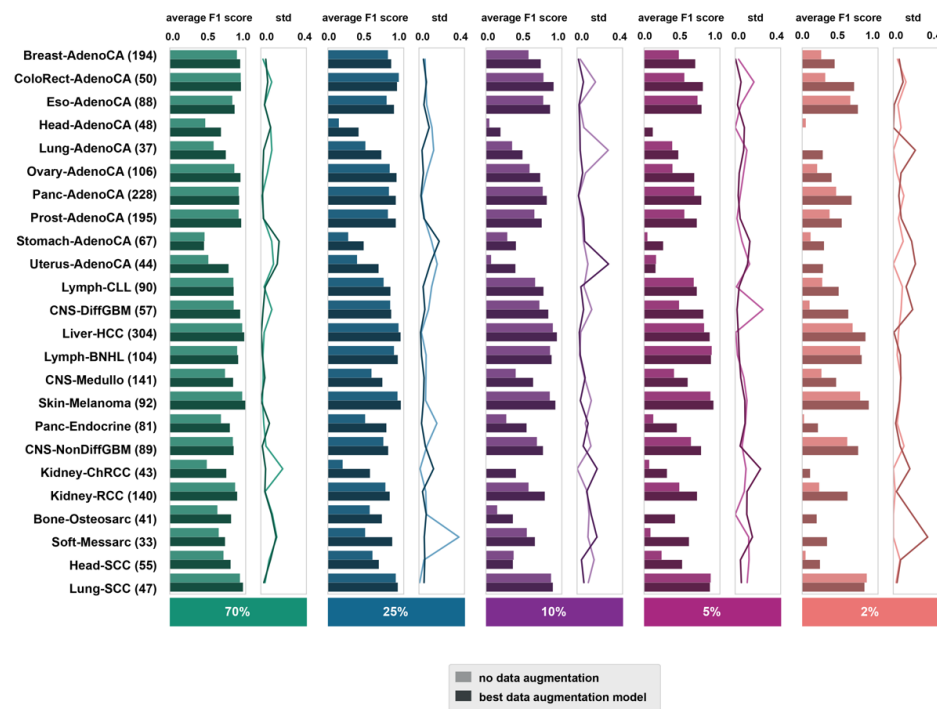
3.3. Data Integration

To compensate for the higher data sparsity, integrating additional genomic features based on the SNV data (i.e., genome-wide SNV density bins, mutation density on driver genes and genome-wide trinucleotide motifs) may improve classification accuracy even further. Model design is a crucial element in this process, even more so when multiple feature types have to be processed in the model. We therefore assessed three main architecture types for the integration of the three SNV-derived features: early integration, multi-branch integration and consecutive integration. The aim of testing these main architecture types was to assess different feature integration approaches, where consecutive integration represents a hierarchical feature integration approach, multi-branch integration applies feature integration in the hidden feature space of the deep learning model, while conversely with early integration we explored feature integration at the raw input feature level. In early integration, the original feed-forward deep learning model was applied where all three feature sets were merged together into one input vector. This architecture is the most simplistic yet it allows the model to extract hidden features using all input features already in the first layer right after the input layer. In multi-branch integration, the feature sets were fed into the model through separated branches and then the final individual dense layers were concatenated before calculating the final output. In consecutive integration, the feature sets are processed one after another and when the processing of a feature set is done, the last dense layer of that feature branch is concatenated with the next feature set into a new input layer (Figure 1d).

For all models, hyperparameter tuning was done by Bayesian optimization in the same fashion as in [5] and the optimization of the number of nodes and layers was performed after each input entry point (Section 2). All data integration experiments were carried out with the best performing data augmentation setting which varied across mutation downsampling rates (70%—20×, 25%—20×, 10%—30×, 5%—50×, 2%—50×).



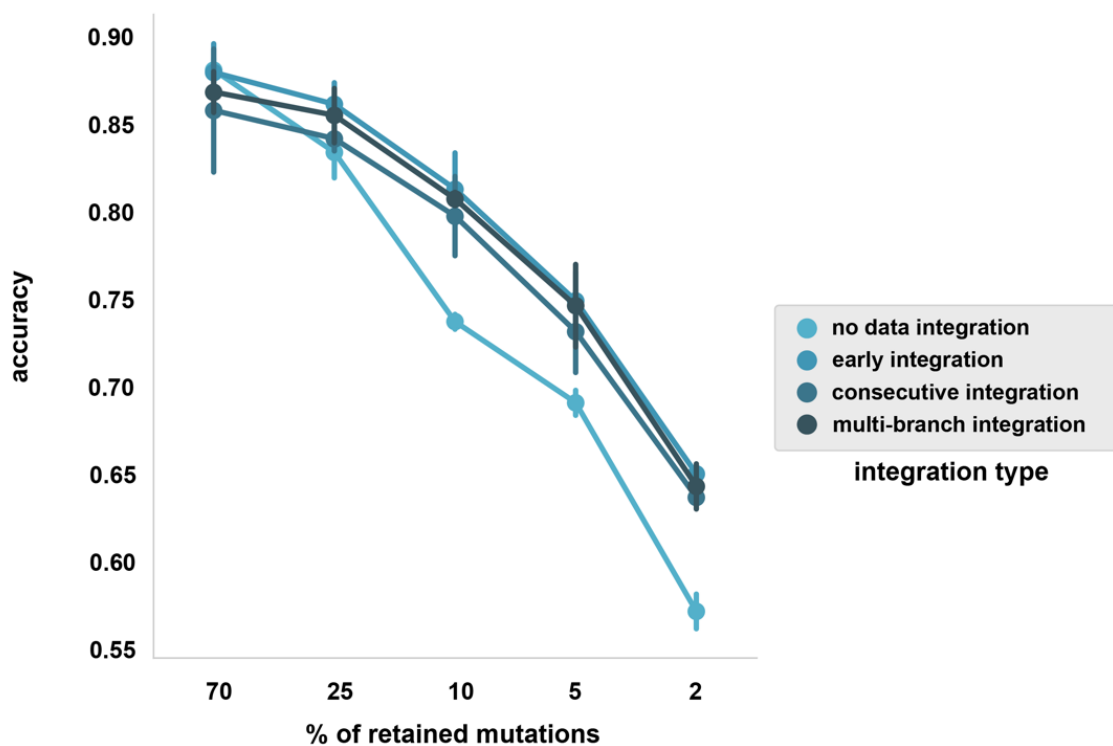
(a)



(b)

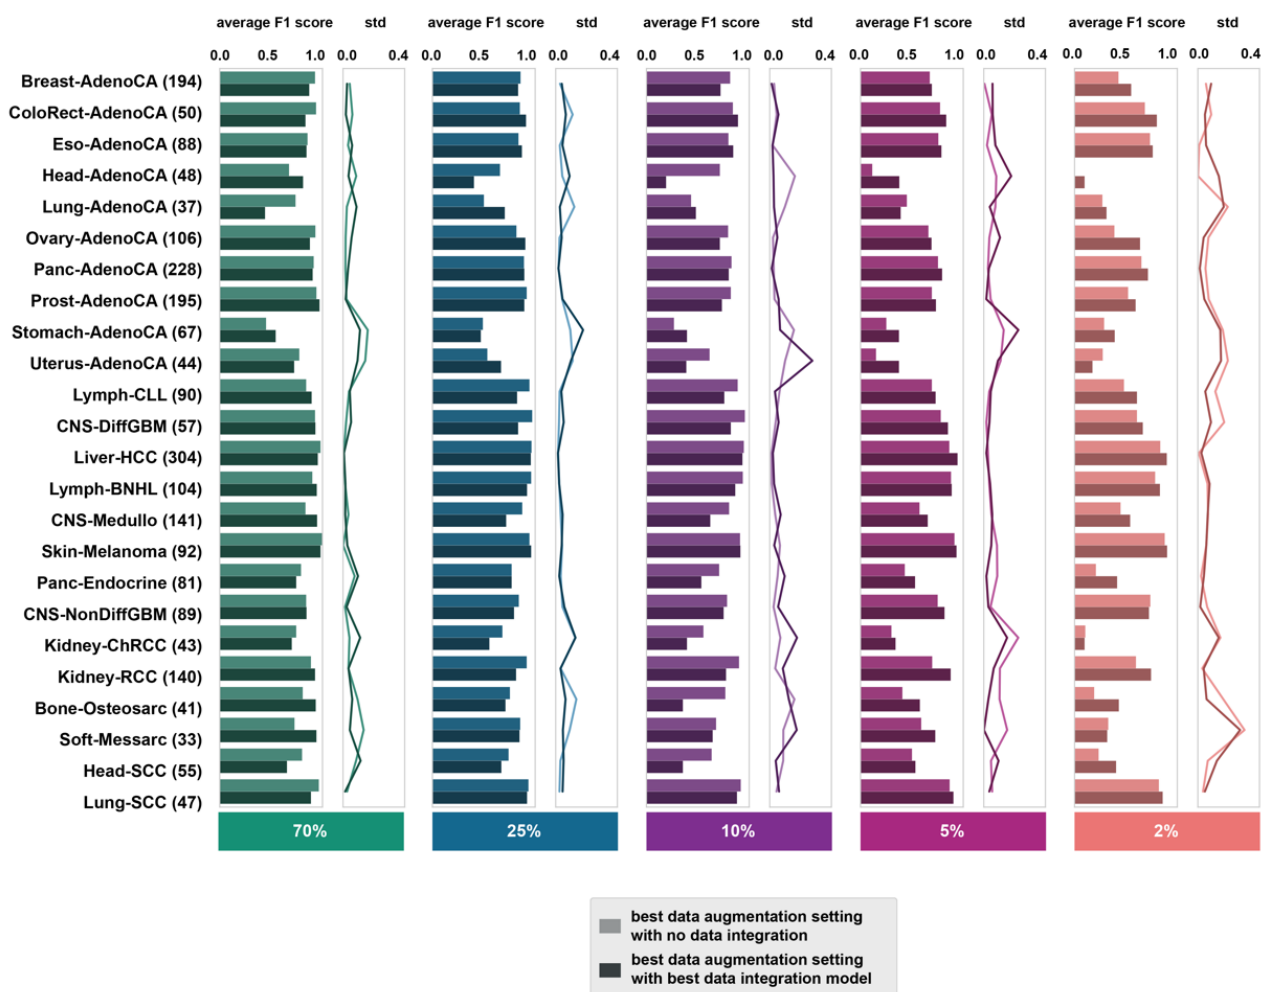
Figure 2. Classifier performance with data augmentation. (a) Overall effect of data augmentation at different mutation downsampling rates (70%, 25%, 10%, 5%, 2%). Average accuracy (% of accurately classified samples, mean value with standard deviation in 3 CV folds) is shown at different downsampling rates with and without applied data augmentation setups (b) Classifier performance at different mutation downsampling rates with the best data augmentation setup, expressed in F1-scores per cancer type (average and std is shown based on 3 folds). The best data augmentation setup was selected based on average accuracy values: 70%—20, 25%—20×, 10%—30×, 5%—50×, 2%—50×.

We found that including the 96 trinucleotide motifs and driver genes can improve the classification accuracy on the sparse somatic data. Average accuracy changes were compared to the best data augmentation models without data integration in paired t-tests where starred comparisons indicate significant differences at $p < 0.05$ (Supplementary Table S5 and S7, Supplementary Figure S5b). Early integration experiments resulted in the highest average accuracies overall, with an average accuracy of 0.65** compared to the data augmented baseline with 0.57 at 2%, with an average accuracy of 0.75* compared to the data augmented baseline with 0.69 at 5%, with an average accuracy of 0.82* compared to the data augmented baseline with 0.74 at 10% and with an average accuracy of 0.86 compared to the data augmented baseline with 0.84 at 25%, with the 70% sparse data experiment being the only exception where the average accuracy remained nearly the same at 0.884 compared to the data augmented baseline with 0.882 (Figure 3a). When assessing the performance of the best data integration model at cancer type level, for most cancer types a gradual accuracy improvement is apparent (Figure 3b). At 2% and 5% downsampling rates, the accuracy of the majority of cancer types shows the most notable improvement, which is especially appreciable in Kidney-RCC, Stomach-AdenoCA, Panc-Endocrine, Bone-Osteosarc and CNS-Medullo (Figure 3b). At other downsampling rates, where more SNVs are retained by default, certain cancer type accuracies slightly decrease (e.g., Stomach-AdenoCA at 10% and Ovary-AdenoCA at 25%). In these cases the data augmentation might be insufficient to combat the increased dimensionality which moreover introduces additional noise caused by the additional features (Figure 3b). We conclude that utilizing additional, compact features derived from the raw SNV data can further aid the cancer type classification model, especially when only 2% or 5% of the somatic SNVs are retained in the sparse somatic samples.



(a)

Figure 3. Cont.



(b)

Figure 3. Classifier performance with data integration. (a) Overall effect of data integration at different mutation downsampling rates (70%, 25%, 10%, 5%, 2%). Average accuracy (% of accurately classified samples, mean value with standard deviation in 3 CV folds) is shown at different downsampling rates with and without different applied data integration architectures (b) Classifier performance at different mutation downsampling rates with the best data integration architecture, expressed in F1-scores per cancer type (average and std is shown based on three folds). The best data integration architecture was selected based on average accuracy (a) values, which was consistently the early data integration architecture across all downsampling rates. For comparison we used the best data augmentation model as the “model without data integration”.

3.4. Assessment of Top 1, 2 and 3 Ranked Predictions

In a diagnostic setting, where the aim of the cancer type classifier is to support the physician in making treatment decisions, it may be sufficient to provide a ranked list of predictions, rather than only the top prediction. This is especially true if a few predictions are close in terms of ranking based on the predicted cancer type probabilities (i.e., have very similar classification scores). These tumor types may have similar features, which could direct treatment decisions as well. We therefore also assessed the classifier accuracy when considering the top 1, 2 and 3 predictions of the model. When comparing the average accuracy of top 1 predictions (average ratio of correctly classified samples in stratified shuffle split cross-validation with three partitions) between different mutation downsampling rates with and without our model improvements, the final improved model achieves an average accuracy of 0.65 compared to an average accuracy of 0.41 with the

original model at 2%, 0.75 compared to 0.56 at 5%, 0.82 compared to 0.63 at 10%, 0.86 compared to 0.75 at 25% and 0.88 compared to 0.83 at 70% (Figure 4a). When the original model trained on full somatic data is compared to different mutation downsampling rates with our model improvements in terms of average accuracy of the top 1 predictions, we found that the classifier performs slightly better at 70% downsampling rate than at 100% (Figure 4a). The underlying reason is likely that the noise introduced by the missing SNV calls functions as a regularizer at this rate and it improves on the generalization performance of the deep learning model during training and validation (at higher downsampling rates, this benefit clearly diminished since at those rates the lack of data is so severe that it decreases the overall performance). When the correct cancer type is identified from the top 3 predictions, the average accuracy is above 0.85 with our improvements, even at 2% downsampling rate (Figure 4a). When considering the top 2 and 3 predictions, the detection of all cancer types is improved, most prominently in case of Stomach-AdenoCA (Figure 4b).

3.5. Feature Importance Assessment

Interpretability is an important aspect in the application of machine learning, especially in the clinical field. By understanding why the model has made a certain prediction, the adoption of the approach in the clinic is facilitated as physicians can, to a certain extent, manually verify its decisions based on their own knowledge and experience. Model interpretation is usually done by feature importance assessment (or in other words feature attribution). However, defining feature importance is inherently problematic with deep learning models, due to the highly non-linear and complex nature of the models. As a result, along with the emergence of these models, solutions to assess feature attributions have also been developed. We aimed to use a straightforward method for our interpretation task and applied the so-called ‘integrated gradients’ method to obtain feature attributions [32] (see Section 2). In summary, the method takes the gradient of the output (probability of cancer type) with respect to the input feature values. In case of positive feature attribution, when the value of the feature was moving away from the baseline value then the probability of the given cancer type would increase, while in case of negative feature attribution, if the feature value moved closer to the baseline value then the target probability would increase [32]. Feature attribution is calculated per sample. To attain a feature importance within the context of a specific cancer type, feature attributions are averaged across the samples within the class.

First, we aimed to assess feature importances on the original, full SNV data and executed early data integration on it, in order to have feature importances calculated based on the full somatic profile of the PCAWG samples. The top 1 bin, top 2 trinucleotide and top 2 driver gene features were selected in each cancer type separately based on the ranked absolute feature importances, and then the union of all top features across all cancer types (41 bins, 13 driver genes and 19 trinucleotides) were visualized (Figure 5a). Interestingly, in Skin-Melanoma, five of the positively attributed top trinucleotide features bear C > T nucleotide transition (CCC > CTC, CCT > CTT, GCA > GTA, TCC > TTC, TCT > TTT), while other positively attributed substitution types only appear in one or two trinucleotides. This is in line with the fact that this cancer type is dominated by C > T nucleotide transitions due to UV-light exposure [38]. *BCL2* has a high positive attribution in Lymph-BNHL and it was found to be frequently mutated in multiple lymphomas [39]. *ERBB4* is among the top attributed driver genes in Liver-HCC, which also corresponds to previous findings, where *ERBB4* was described as a suppressor in the development of Liver-HCC, suggesting that somatic mutations in *ERBB4* can function as drivers [40]. *ERBB4* also positively contributed to the prediction of Stomach-AdenoCA, Panc-AdenoCA, Eso-AdenoCA, Kidney-RCC and Lymph-CLL. *ERBB4* mutations were found in gastric carcinomas, albeit to a lesser extent than in colorectal and lung carcinomas [41], however, to our knowledge, notable SNVs have not been reported in Lymph-CLL. Moreover, *ERBB4* was proposed as a tumor suppressor in Kidney-RCC [42], which might correlate with its positive feature attribution driven by increased somatic mutation density in *ERBB4* in Kidney-RCC. Importantly, we do not

imply to draw direct connections with causality but rather aim to show that certain feature importances can capture meaningful underlying biological variation.

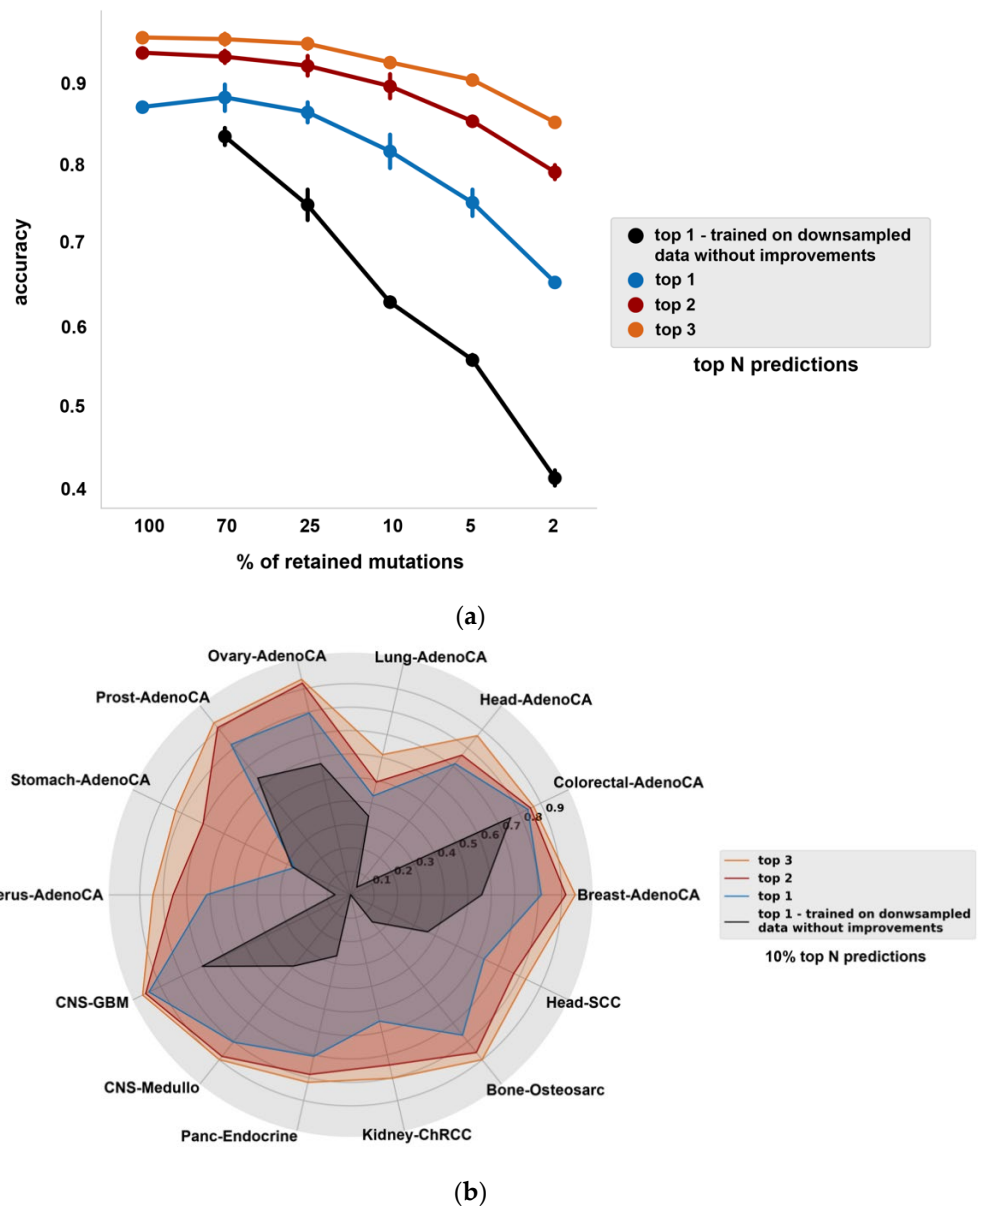


Figure 4. Classifier performance in predicting the top N labels. **(a)** Average accuracy (% of accurately classified samples, mean value with standard deviation in 3 CV folds) when the correct cancer type is among top 1, 2 and 3 predictions. Average accuracy of the best models is shown at each mutation downsampling rate when identifying the correct tumour type among its top N -ranked predictions $N \in (1, 2, 3)$, error bars indicate standard deviation. **(b)** Classifier performance on somatic SNV data with 10% of retained mutations when the correct cancer type is among top 1, 2 and 3 predictions, compared to the baseline model. The axes represent the accuracy for each cancer type, expressed in average F1 scores (stratified shuffle split cross-validation with three partitions).

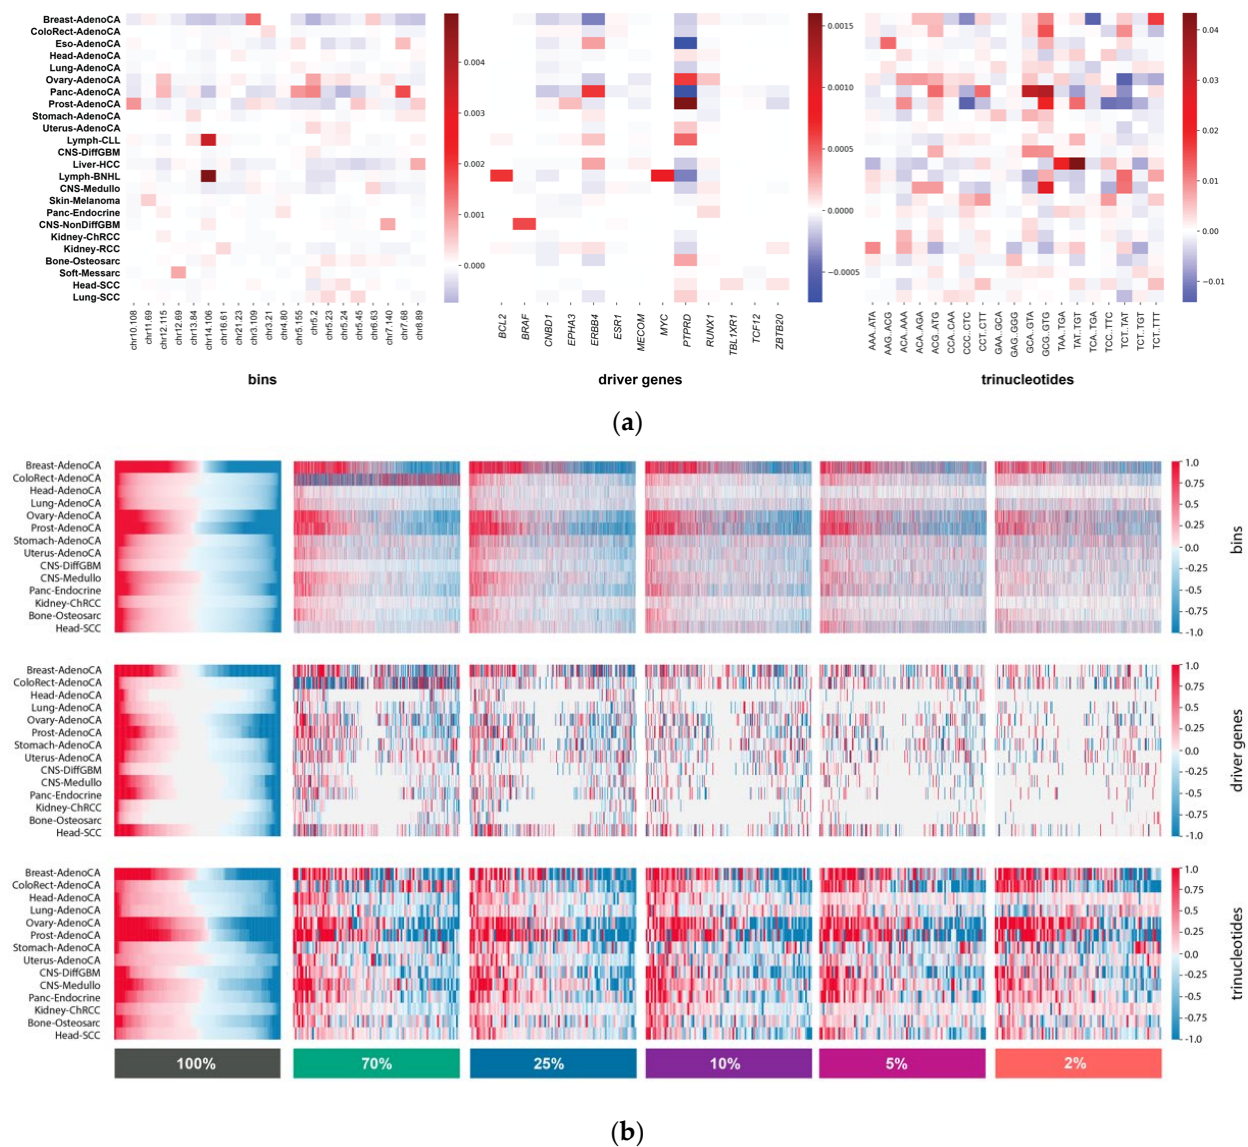


Figure 5. Feature importance assessment across all sparsity levels. (a) Top features in each feature type based on feature importance in full solid biopsy data. Top features were selected based on average feature importance ranking per cancer type in the test folds, resulting in 41 bins, 13 driver genes and 19 trinucleotides. The color scale represents negative (blue) and positive (red) feature importance values with zero (white) centering. (b) Feature importance patterns per feature type in full solid biopsy SNV data and sparse SNV data. Heatmaps of overall feature importance patterns were log-modulus transformed and scaled. Feature attributions were sorted per cancer type in the full solid biopsy SNV data and then the sparse data matrices were sorted to follow the same feature order per cancer type.

Overall changes in feature importance were compared between the original, full somatic data and sparse data settings in early data integration, as this architecture type resulted in the highest accuracy improvements in general. Our results show that the feature importance changes considerably in sparse data compared to the full SNV data (Figure 5b). The effect of SNV downsampling can be clearly seen when comparing feature importance profiles of the bin features in the baseline experiment with full somatic data to the experiments with sparse data, where many features have changed signs or have lessened or completely diminished importance due to sparsity. On the other hand, the 96 trinucleotide feature importances did not change drastically in sparse data, indicating that

these features can still be harnessed more effectively by the classifier. This again validates our intuition, as we expected to see performance improvement on much sparser data with the inclusion of more SNV derived features.

4. Discussion

Diagnosis of CUPs poses an important challenge in clinical practice. Utilizing the genome-wide somatic mutation profiles of ctDNA could offer a valuable approach to detect the cell of origin of cancer, but liquid biopsies bring about their own challenges as the inherently low concentration of ctDNA gives rise to sparse somatic mutation profiles. Our results show that both data augmentation and data integration can massively improve classifier performance, but the effect depends on the sparsity of mutational profiles. It appears that at less severe 70% and 25% downsampling, where more information is preserved from the original WGS samples by definition, the classifier cannot exploit the generated training samples of higher data augmentation levels (30–50×) and, hence, 20× data augmentation is sufficient to improve classifier performance. Data integration resulted in performance improvement in all sparse datasets except 70%, where the average accuracy remained nearly the same in comparison to the best data augmented model (0.884 and 0.882). When looking at the top 1 predictions with our improved model compared to the predictions made with the original model in all sparse datasets, the improved model achieves an average accuracy of 0.65 compared to an average accuracy of 0.41 with the original model at 2%, 0.75 compared to 0.56 at 5%, 0.82 compared to 0.63 at 10%, 0.86 compared to 0.75 at 25% and 0.88 compared to 0.83 at 70%, respectively, while in case the correct cancer type is identified from the top 3 predictions, our model reaches an average accuracy above 0.85, even at 2% and 5% downsampling rate (Figure 4a). This suggests that in case targeted treatments can be sufficiently selected based on the top 3 predicted cancer types, our classifier can aid treatment decisions even in extremely sparse data conditions.

In previous studies, the reported percentage of somatic mutations that can be detected in liquid biopsies compared to solid tissue biopsies in most advanced and metastatic cancers were 72% and 73.8% [17,19], therefore our 70% sparse dataset was designated as an approximation of the reported values that can potentially reflect the sparsity in metastatic/CUP liquid biopsy data. As mentioned above, the final average accuracy at 70% is 0.88 compared to the baseline 0.83, and might indicate that this genome-wide SNV mutation information based approach can perform well in advanced and metastatic cancers. Interestingly, at 70%, the average accuracy has improved compared to the baseline model with full somatic SNV data (Figure 4a), which suggests that mutation downsampling can function as a regularizer at this rate, giving rise to improved performance. However, while retaining 70% of somatic SNVs provides a reasonable dataset to get an accuracy estimate for most tumor types in CUP/metastatic condition, this approximation might only be reliably applied for samples which are deep sequenced and cannot hold for samples with lower average coverage. Besides, certain cancer types such as glioma, medulloblastoma, bladder, renal and gastroesophageal cancer have been reported to have significantly lower ctDNA concentrations even in advanced disease [10]. For these tumor types, fewer variants will be detected in the blood and, therefore, lower downsampling rates provide a better estimation of classifier performance. With the improvements applied on the sparse dataset, a classification accuracy of 0.87 was achieved at 25% and 0.8 at 10% in CNS-Medullo (medulloblastoma), 0.97 accuracy was achieved at 25% and 0.95 at 10% in CNS-DiffGBM (diffuse glioma), 0.84 accuracy was achieved at 25% and 0.78 at 10% in CNS-DiffGBM (non-diffuse glioma), 0.68 accuracy was achieved at 25% and 0.55 at 10% in Kidney-ChRCC (renal cell carcinoma, distal tubule), 0.92 accuracy was achieved at 25% and 0.9 at 10% in Kidney-RCC (renal cell carcinoma, proximal tubule) and a classification accuracy of 0.84 was achieved at 25% and 0.79 at 10% in Eso-AdenoCA (esophageal adenocarcinoma).

The results presented here indicate that most of the problematic cancer types can still be classified with a reasonable accuracy even in very sparse somatic SNV data, which supports the diagnostic relevance of somatic SNV data derived from liquid biopsies. Moreover, since

variants at low allele frequencies can be missed in solid biopsy samples at lower tumor purity levels, we hypothesise the applicability of our improvements to boost predictive performance on tissue biopsy samples, by means of predicting the cell of origin in tissue biopsy data with a classification model that was trained on sparse augmented data with data integration. This argument is bolstered by our finding where we compared the performance of the improved classifier on 70% downsampled data to the performance of the original model on the 100% SNV data, and the average accuracy of the 70% model was found to be improved compared to the baseline model (Figure 4a).

An important caveat in using somatic mutation profiles obtained from liquid biopsies is the presence of non-cancerous somatic mutations. According to a recent study, most of the non-cancerous cfDNA originates from white blood cells (WBC), therefore liquid biopsies contain a significant amount of somatic mutations as a result of clonal haematopoiesis in WBC [6,17], which are indistinguishable from somatic mutations that arise in the cancer. This finding emphasises the importance of matched cfDNA WBC sequencing for accurate somatic variant interpretation in liquid biopsies. In our work, we only accounted for the sparse profile of somatic mutations that arise from ctDNA and considered the non-cancerous somatic SNVs and germline mutations to be already filtered out based on e.g., matched cfDNA WBC sequencing.

5. Conclusions

Despite the challenge of using sparse somatic data to identify the cell of origin of cancer, our results demonstrate that a deep learning classification model that uses multiple feature types and training set augmentation can help to harness the available sparse data in an efficient way. We conclude that somatic mutation density profiles obtained from liquid biopsies are suitable for the detection of the cell of origin of cancer, in particular for advanced and metastatic patients, and therefore they provide a valuable alternative to invasive solid tissue biopsies for CUP identification.

Supplementary Materials: The following are available online at <https://www.mdpi.com/article/10.3390/life12010001/s1>, Table S1: Sample distribution across cancer types in the used PCAWG dataset, Table S2: List of used PCAWG samples in the training, validation and test folds with donor identifiers and DCC project codes, Table S3: List of final optimized hyperparameters of all models, Table S4.: Number of training samples with and without data augmentation in the 3 cross-validation folds, Table S5: Saphiro-Wilk test of normal distribution in each data augmentation and data integration experiment. Available data points (accuracies in 3-fold CV) were limited, Table S6: Paired t-test statistics and p-value in data augmentation experiments. All data augmentation experiments were compared to the baseline 3-fold CV accuracies (model without data augmentation), Table S7: Paired t-test statistics and p-value in data integration experiments. All data integration experiments were compared to the best data augmentation model 3-fold CV accuracies (without any data integration), Figure S1: Classifier performance at different mutation downsampling rates (expressed in average F1 score and standard deviation of accuracy in folds) with the original model trained on the full (100%) data. Test samples were generated by retaining 70%, 25%, 10%, 5% and 2% of somatic mutations from each original sample (mutation downsampling) while the training and validation set consisted of only original, full somatic SNV samples, Figure S2: Classifier performance (expressed in average F1 score and standard deviation of accuracy in folds) at different mutation downsampling rates with models trained on sparse data sets. Samples were generated by retaining 70%, 25%, 10%, 5% and 2% of somatic mutations from each original sample (mutation downsampling) for the training, validation and test sets, Figure S3: Classifier performance (expressed in average F1 score and standard deviation of accuracy in folds) at different mutation downsampling rates with all the data augmentation setups. Samples were generated by retaining 70%, 25%, 10%, 5% and 2% of somatic mutations from each original sample (mutation downsampling). Different data augmentation levels were assessed, see Data augmentation in Materials and methods, Figure S4: Classifier performance (expressed in average F1 score and standard deviation of accuracy in folds) at different mutation downsampling rates with all the data integration setups on the best data augmentation model. Samples were generated by retaining 70%, 25%, 10%, 5% and 2% of somatic mutations from each original sample (mutation downsampling). After performing data augmentation

experiments, the best data augmentation setup was selected based on average accuracy values: 70%—20×, 25%—20×, 10%—30×, 5%—50×, 2%—50×. Figure S5. Categorical heatmap of p-values in paired t-tests of data augmentation experiments (compared to baseline) and data integration experiments (compared to best data augmentation model).

Author Contributions: Conceptualization, J.d.R. and A.D.; methodology, J.d.R. and A.D.; software, A.D.; validation, A.D.; formal analysis, A.D.; investigation, M.J. and A.D.; resources, J.d.R.; data curation, A.D.; writing—original draft preparation, J.d.R., M.J. and A.D.; writing—review and editing, J.d.R., M.J. and A.D.; visualization, A.D.; supervision, J.d.R.; project administration, J.d.R.; funding acquisition, J.d.R. All authors have read and agreed to the published version of the manuscript.

Funding: J.d.R. and A.D. are supported by a Vidi Fellowship (639.072.715) from the Dutch Organization for Scientific Research (Nederlandse Organisatie voor Wetenschappelijk Onderzoek, NWO).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The used PCAWG (<https://doi.org/10.1038/s41586-020-1969-6>, accessed on 16 January 2018) sample identifiers are listed in Supplementary Table S2. Aligned sequencing data, somatic and germline variant calls are available at <https://dcc.icgc.org/releases/PCAWG>, accessed on 16 January 2018. Further information regarding data access can be found at <https://docs.icgc.org/pcawg/data/>, accessed on 16 January 2018. Scripts that were used for classification, feature importance calculation and sparse augmented data generation are available at https://github.com/UMCUGenetics/cancer_type_classification_from_sparse_SNV_data, accessed on 1 March 2020, on trained models are also available upon request. The code is distributed under the Apache Version 2.0 Open Source license (<https://www.apache.org/licenses/LICENSE-2.0>, accessed on 1 March 2020).

Conflicts of Interest: J.d.R. co-founded Cyclomics and M.J. is an academic collaborator at Cyclomics. A.D. declares no conflict of interests. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

References

1. Goldie, S.J.; Chincarini, G.; Darido, C. Targeted Therapy Against the Cell of Origin in Cutaneous Squamous Cell Carcinoma. *Int. J. Mol. Sci.* **2019**, *20*, 2201. [[CrossRef](#)] [[PubMed](#)]
2. Greco, F.A. Molecular diagnosis of the tissue of origin in cancer of unknown primary site: Useful in patient management. *Curr. Treat. Options Oncol.* **2013**, *14*, 634–642. [[CrossRef](#)] [[PubMed](#)]
3. Pavlidis, N.; Khaled, H.; Gaafar, R. A mini review on cancer of unknown primary site: A clinical puzzle for the oncologists. *J. Adv. Res.* **2015**, *6*, 375–382. [[CrossRef](#)]
4. Salvadores, M.; Mas-Ponte, D.; Supek, F. Passenger mutations accurately classify human tumors. *PLoS Comput. Biol.* **2019**, *15*, e1006953. [[CrossRef](#)]
5. Jiao, W.; Atwal, G.; Polak, P.; Karlic, R.; Cuppen, E.; Danyi, A.; de Ridder, J.; van Herpen, C.; Lolkema, M.P.; Steeghs, N.; et al. A deep learning system accurately classifies primary and metastatic cancers using passenger mutation patterns. *Nat. Commun.* **2020**, *11*, 728. [[CrossRef](#)]
6. Overman, M.J.; Modak, J.; Kopetz, S.; Murthy, R.; Yao, J.C.; Hicks, M.E.; Abbruzzese, J.L.; Tam, A.L. Use of research biopsies in clinical trials: Are risks and benefits adequately discussed? *J. Clin. Oncol.* **2013**, *31*, 17–22. [[CrossRef](#)]
7. Cohen, J.D.; Li, L.; Wang, Y.; Thoburn, C.; Afsari, B.; Danilova, L.; Douville, C.; Javed, A.A.; Wong, F.; Mattox, A.; et al. Detection and localization of surgically resectable cancers with a multi-analyte blood test. *Science* **2018**, *359*, 926–930. [[CrossRef](#)] [[PubMed](#)]
8. Sung, J.S.; Chong, H.Y.; Kwon, N.J.; Kim, H.M.; Lee, J.W.; Kim, B.; Lee, S.B.; Park, C.W.; Choi, J.Y.; Chang, W.J.; et al. Detection of somatic variants and *EGFR* mutations in cell-free DNA from non-small cell lung cancer patients by ultra-deep sequencing using the ion ampliseq cancer hotspot panel and droplet digital polymerase chain reaction. *Oncotarget* **2017**, *8*, 106901–106912. [[CrossRef](#)] [[PubMed](#)]
9. Iwahashi, N.; Sakai, K.; Noguchi, T.; Yahata, T.; Matsukawa, H.; Toujima, S.; Nishio, K.; Ino, K. Liquid biopsy-based comprehensive gene mutation profiling for gynecological cancer using CAnCER Personalized Profiling by deep Sequencing. *Sci. Rep.* **2019**, *9*, 10426. [[CrossRef](#)] [[PubMed](#)]
10. Bettgowda, C.; Sausen, M.; Leary, R.J.; Kinde, I.; Wang, Y.; Agrawal, N.; Bartlett, B.R.; Wang, H.; Luber, B.; Alani, R.M.; et al. Detection of circulating tumor DNA in early- and late-stage human malignancies. *Sci. Transl. Med.* **2014**, *6*, 224ra24. [[CrossRef](#)] [[PubMed](#)]

11. Haber, D.A.; Velculescu, V.E. Blood-based analyses of cancer: Circulating tumor cells and circulating tumor DNA. *Cancer Discov.* **2014**, *4*, 650–661. [[CrossRef](#)]
12. Mouliere, F.; Smith, C.G.; Heider, K.; Su, J.; van der Pol, Y.; Thompson, M.; Morris, J.; Wan, J.C.M.; Chandrananda, D.; Hadfield, J.; et al. Fragmentation patterns and personalized sequencing of cell-free DNA in urine and plasma of glioma patients. *EMBO Mol. Med.* **2021**, *13*, e12881. [[CrossRef](#)]
13. Ma, X.; Zhu, L.; Wu, X.; Bao, H.; Wang, X.; Chang, Z.; Shao, Y.W.; Wang, Z. Cell-Free DNA Provides a Good Representation of the Tumor Genome Despite Its Biased Fragmentation Patterns. *PLoS ONE* **2017**, *12*, e0169231. [[CrossRef](#)] [[PubMed](#)]
14. Cristiano, S.; Leal, A.; Phallen, J.; Fiksel, J.; Adleff, V.; Bruhm, D.C.; Jensen, S.Ø.; Medina, J.E.; Hruban, C.; White, J.R.; et al. Genome-wide cell-free DNA fragmentation in patients with cancer. *Nature* **2019**, *570*, 385–389. [[CrossRef](#)]
15. Keup, C.; Suryaprakash, V.; Hauch, S.; Storbeck, M.; Hahn, P.; Sprenger-Haussels, M.; Kolberg, H.C.; Tewes, M.; Hoffmann, O.; Kimmig, R.; et al. Integrative statistical analyses of multiple liquid biopsy analytes in metastatic breast cancer. *Genome Med.* **2021**, *13*, 85. [[CrossRef](#)] [[PubMed](#)]
16. Higgins, M.J.; Jelovac, D.; Barnathan, E.; Blair, B.; Slater, S.; Powers, P.; Zorzi, J.; Jeter, S.C.; Oliver, G.R.; Fetting, J.; et al. Detection of tumor PIK3CA status in metastatic breast cancer using peripheral blood. *Clin. Cancer Res.* **2012**, *18*, 3462–3469. [[CrossRef](#)] [[PubMed](#)]
17. Razavi, P.; Li, B.T.; Brown, D.N.; Jung, B.; Hubbell, E.; Shen, R.; Abida, W.; Juluru, K.; De Bruijn, I.; Hou, C.; et al. High-intensity sequencing reveals the sources of plasma circulating cell-free DNA variants. *Nat. Med.* **2019**, *25*, 1928–1937. [[CrossRef](#)]
18. Adalsteinsson, V.A.; Ha, G.; Freeman, S.S.; Choudhury, A.D.; Stover, D.G.; Parsons, H.A.; Gydush, G.; Reed, S.C.; Rotem, D.; Rhoades, J.; et al. Scalable whole-exome sequencing of cell-free DNA reveals high concordance with metastatic tumors. *Nat. Commun.* **2017**, *8*, 1324. [[CrossRef](#)] [[PubMed](#)]
19. Jiang, J.; Adams, H.P.; Yao, L.; Yaung, S.; Lal, P.; Balasubramanyam, A.; Fuhlbrück, F.; Tikoo, N.; Lovejoy, A.F.; Froehler, S.; et al. Concordance of Genomic Alterations by Next-Generation Sequencing in Tumor Tissue versus Cell-Free DNA in Stage I–IV Non-Small Cell Lung Cancer. *J. Mol. Diagn.* **2020**, *22*, 228–235. [[CrossRef](#)]
20. McCabe, M.J.; Gauthier, M.E.A.; Chan, C.L.; Thompson, T.J.; De Sousa, S.M.; Puttick, C.; Grady, J.P.; Gayevskiy, V.; Tao, J.; Ying, K.; et al. Development and validation of a targeted gene sequencing panel for application to disparate cancers. *Sci. Rep.* **2019**, *9*, 17052. [[CrossRef](#)] [[PubMed](#)]
21. Zviran, A.; Schulman, R.C.; Shah, M.; Hill, S.T.K.; Deochand, S.; Khamnei, C.C.; Maloney, D.; Patel, K.; Liao, W.; Widman, A.J.; et al. Genome-wide cell-free DNA mutational integration enables ultra-sensitive cancer monitoring. *Nat. Med.* **2020**, *26*, 1114–1124. [[CrossRef](#)] [[PubMed](#)]
22. The ICGC/TCGA Pan-Cancer Analysis of Whole Genomes Consortium; Campbell, P.J.; Getz, G. Pan-cancer analysis of whole genomes. *Nature* **2020**, *578*, 82–93. [[CrossRef](#)]
23. Alexandrov, L.B.; Nik-Zainal, S.; Wedge, D.C.; Aparicio, S.A.; Behjati, S.; Biankin, A.V.; Bignell, G.R.; Bolli, N.; Borg, A.; Børresen-Dale, A.L.; et al. Signatures of mutational processes in human cancer. *Nature* **2013**, *500*, 415–421. [[CrossRef](#)] [[PubMed](#)]
24. Hanahan, D.; Weinberg, R.A. Hallmarks of cancer: The next generation. *Cell* **2011**, *144*, 646–674. [[CrossRef](#)]
25. Schaefer, M.; Serrano, L. Cell type-specific properties and environment shape tissue specificity of cancer genes. *Sci. Rep.* **2016**, *6*, 20707. [[CrossRef](#)] [[PubMed](#)]
26. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. *arXiv* **2014**, arXiv:1412.6980v9.
27. Head, T.; Kumar, M.; Nahrstaedt, H.; Louppe, G.; Shcherbatyi, I. Scikit-Optimize/Scikit-Optimize (v0.7.4). *Zenodo* **2020**. [[CrossRef](#)]
28. Glorot, X.; Bengio, Y. Understanding the difficulty of training deep feedforward neural networks. In Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, PMLR, Sardinia, Italy, 13–15 May 2010; Volume 9, pp. 249–256. Available online: <http://proceedings.mlr.press/v9/glorot10a.html> (accessed on 1 August 2020).
29. Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. Scikit-Learn: Machine Learning in Python. *JMLR* **2011**, *12*, 2825–2830. Available online: <https://jmlr.csail.mit.edu/papers/v12/pedregosa11a.html> (accessed on 1 January 2020).
30. Bailey, M.H.; Tokheim, C.; Porta-Pardo, E.; Sengupta, S.; Bertrand, D.; Weerasinghe, A.; Colaprico, A.; Wendl, M.C.; Kim, J.; Reardon, B.; et al. Comprehensive Characterization of Cancer Driver Genes and Mutations. *Cell* **2018**, *173*, 371–385.e18. [[CrossRef](#)]
31. Sundararajan, M.; Taly, A.; Yan, Q. Axiomatic Attribution for Deep Networks. *arXiv* **2017**, arXiv:1703.01365v2.
32. Shorten, C.; Khoshgoftaar, T.M. A survey on Image Data Augmentation for Deep Learning. *J. Big Data* **2019**, *6*, 60. [[CrossRef](#)]
33. Zhong, Z.; Liang, Z.; Guoliang, K.; Shaozi, L.; Yi, Y. Random Erasing Data Augmentation. *arXiv* **2020**, arXiv:1708.04896v2. [[CrossRef](#)]
34. Krizhevsky, A.; Sutskever, I.; Hinton, G.E. ImageNet classification with deep convolutional neural networks. *Commun. ACM* **2017**, *60*, 84–90. [[CrossRef](#)]
35. Summers, C.; Dinneen, M.J. Improved mixed-example data augmentation. *arXiv* **2018**, arXiv:1805.11272v4.
36. Inoue, H. Data Augmentation by Pairing Samples for Images Classification. *arXiv* **2018**, arXiv:1801.02929v2.
37. Hestness, J.; Narang, S.; Ardalani, N.; Diamos, G.; Jun, H.; Kianinejad, H.; Patwary, M.; Ali, M.; Yang, Y.; Zhou, Y. Deep Learning Scaling is Predictable, Empirically. *arXiv* **2017**, arXiv:1712.00409v1.
38. Hodis, E.; Watson, I.R.; Kryukov, G.V.; Arold, S.T.; Imielinski, M.; Theurillat, J.P.; Nickerson, E.; Auclair, D.; Li, L.; Place, C.; et al. A landscape of driver mutations in melanoma. *Cell* **2012**, *150*, 251–263. [[CrossRef](#)]

39. Schuetz, J.M.; Johnson, N.A.; Morin, R.D.; Scott, D.W.; Tan, K.; Ben-Nierah, S.; Boyle, M.; Slack, G.W.; Marra, M.A.; Connors, J.M.; et al. BCL2 mutations in diffuse large B-cell lymphoma. *Leukemia* **2012**, *26*, 1383–1390. [[CrossRef](#)]
40. Liu, Y.; Song, L.; Ni, H.; Sun, L.; Jiao, W.; Chen, L.; Zhou, Q.; Shen, T.; Cui, H.; Gao, T.; et al. ERBB4 acts as a suppressor in the development of hepatocellular carcinoma. *Carcinogenesis* **2017**, *38*, 465–473. [[CrossRef](#)] [[PubMed](#)]
41. Soung, Y.H.; Lee, J.W.; Kim, S.Y.; Wang, Y.P.; Jo, K.H.; Moon, S.W.; Park, W.S.; Nam, S.W.; Lee, J.Y.; Yoo, N.J.; et al. Somatic mutations of the ERBB4 kinase domain in human cancers. *Int. J. Cancer* **2006**, *118*, 1426–1429. [[CrossRef](#)]
42. Thomasson, M.; Hedman, H.; Junttila, T.T.; Elenius, K.; Ljungberg, B.; Henriksson, R. ErbB4 is downregulated in renal cell carcinoma—A quantitative RT-PCR and immunohistochemical analysis of the epidermal growth factor receptor family. *Acta Oncol.* **2004**, *43*, 453–459. [[CrossRef](#)] [[PubMed](#)]