

Universal Laboratory Model: prognosis of abnormal clinical outcomes based on routine tests

Pavel Karpov
Department of medicinal chemistry
Moscow State University
Moscow, Russia
carpovpv @ qsar.chem.msu.ru

Ilya Petrenkov
Department of computer science
Samara National Research University
Samara, Russia
petilia2002 @ yandex.ru

Ruslan Raiman
R&D in medical laboratories
Ltd, RosLIS
Moscow, Russia
rulsan.raiman @ roslis.ru

Abstract—Clinical laboratory results are ubiquitous in any diagnosis making. Predicting abnormal values of not prescribed tests based on the results of performed tests looks intriguing, as it would be possible to make early diagnosis available to everyone. The special place is taken by the Common Blood Count (CBC) test, as it is the most widely used clinical procedure. Combining routine biochemical panels with CBC presents a set of test-value pairs that varies from patient to patient, or, in common settings, a table with missing values. Here we formulate a tabular modeling problem as a set translation problem where the source set comprises pairs of GPT-like label column embedding and its corresponding value while the target set consists of the same type embeddings only. The proposed approach can effectively deal with missing values without implicitly estimating them and bridges the world of LLM with the tabular domain. Applying this method to clinical laboratory data, we achieve an improvement up to 8% AUC for joint predictions of high uric acid, glucose, cholesterol, and low ferritin levels.

Index Terms—pathology prediction, tabular data, CBC blood test, anemia, uric acid, Set Transformer, Universal Laboratory Model, ULM, RosLIS, Laboratory Information System, LIS agent, missing values

I. INTRODUCTION

The World Health Organization places the complete blood count (CBC) test in the list of basic laboratory tests [1], making it one of the most widely used diagnostic procedures. As the CBC test could be treated as a proxy [2], [3] to determine the pathological states of a patient, modeling these states based on the results of the CBC and other laboratory tests [4]–[11] could early diagnose chronic diseases, preventing them from developing and improving overall well-being at the individual and state levels. Although much effort is being made to develop large language models to analyze textual information in medical documentation [12], the use of ordinary tabular data with laboratory results is less common due to the lack of data and regular difficulties in tabular modeling [13]. In this work, we focus on modeling abnormal values of one test (namely: ferritin [4], [5], uric acid [9]–[11], cholesterol and glucose [7])¹ based solely on demographic data (gender

and age) and the results of other tests (CBC and biochemical tests, if available). Thus, if a patient has CBC results, with the help of the presented approach and models, she could also have a prognosis of low ferritin levels (anemia), high uric acid (gout), high cholesterol (risk of cardiovascular disorders), and high glucose (diabetes).

Most studies explore relatively small datasets from one hospital and use a fixed set of tests (features, indicators) for each sample (classic tabular data) to develop predictive models. As a result, these datasets often have a predominant proportion of sick patients, resulting in remarkably good statistics, Table I.

Low ferritin levels are predicted in [4] to detect iron deficiency anemia (IDA) using random forest-based algorithms (RF), achieving the area under the ROC curve (AUC) of 0.92. Reference [6] reports the application of the gradient boost method (GB) to detect IDA and the best model reaches the AUC of 0.87, for their initial dataset, and even higher 0.89 for the so-called Kenyan dataset.

Detection of the high level of glycated hemoglobin (HbA1c) in [7] with the help of an artificial neural network (ANN) achieves 78.1% sensitivity, 78.7% precision, and an F1 score of 78.4% in diabetes diagnosis. In [8], the high-level HbA1c level is modeled, resulting in an AUC of 0.81.

Models for predicting high uric acid values are presented in [9], where the stacking ensemble model (SEM) reveals the best results: 0.91 accuracy, 0.89 sensitivity, 0.91 specificity, and an F1 score of 0.65. The authors [14] propose a regression model for uric acid, achieving a determination coefficient of 0.99.

The use of fixed sets of tests to build a model contradicts common clinical practice, where a set of prescribed tests for one patient would likely be different from the set of tests for another patient as the choice is usually dependent on the doctors, patients' jobs, individual health state, etc. To improve the accuracy of the predictions, it is important to consider all available information. Therefore, the input data for the predictive models in this case is a sparse matrix, Table III. This leads to the problem of "missing values", where different patients lack different tests, complicating the application of

This work was sponsored by Ltd RosLIS, Russia, <https://roslis.ru>

¹The choice of tests is dictated by the data availability.

Table I: An overview of laboratory data modeling

Test	Data size	Thresholds ^a	AUC ^b	Methods	Refs
Glycated hemoglobin, HbA1c	45 431	$\geq 6.5\%$	0.81	AdaBoost, ANN, LogitBoost, RF, RUSBoost	[8]
Hemoglobin, (HGB)	364	≤ 130 g/L for men (M), ≤ 120 g/L for women (F)	0.97	ANN, Decision Tree (DT), Logistic Regression (LR), NaïveBayes (NB), RF, Support Vector Machine (SVM)	[15]
Ferritin, (FER)	12 009	$\leq 10 - 13$ $\mu\text{g/L}$ for women, $\leq 22 - 30$ $\mu\text{g/L}$ for men ^c	0.92	RF	[4]
Ferritin, (FER)	5 128	≤ 10 $\mu\text{g/L}$ for women ≤ 30 $\mu\text{g/L}$ for men	0.97	Bayesian linear regression, Lasso regression, Linear regression, RF	[5]
Iron deficiency anemia, (IDA)	19 000	≤ 15 $\mu\text{g/L}$	0.87	GB, K-Nearest Neighbors (KNN), LR, NB, RF, XGBoost (XGB)	[6]
Uric Acid, (URIC)	40 899	$\geq 0, 42$ mmol/L for men $\geq 0, 36$ mmol/L for women	0.85	DT, SEM, SVM, XGB	[9]
Uric Acid, (URIC)	91 690	—"—"	0.82	LR, RF	[11]
Uric Acid, (URIC)	38 001	—"—"	0.78	Discrimination analysis classification, DT, KNN, NB, RF, SVM	[10]

^aIf the condition is true, the point belongs to the positive class. ^bBest AUC values are reported if there are many methods used in the article. ^c The authors used different thresholds depending on the laboratory equipment.

standard machine learning methods.

There are many approaches to handling missing values, including machine learning-based methods [16], mean imputation [17], genetic algorithms [18], [19], Generative Adversarial Networks (GANs) [20], [21], etc.

For example, in [5], the authors restore the missing values before predicting low ferritin levels. The four imputation techniques were used: mean, multiple imputation with chained equations–full (MICE–full), multiple imputation with chained equations–select (MICE–sel) and missForest. In the articles [15], [22], the authors predict the missing hemoglobin values (make a regression) before determining the type of anemia.

All imputation methods rely on assumptions about the data distribution, inevitably leading to distortions. Laboratory results are individual for each patient. This uniqueness should be carefully considered when developing machine learning models. A more effective approach may be to adapt the predictive model to handle missing values rather than trying to fill them in.

Our contributions in the article are as follows:

- We developed a new neural network architecture – the Universal Laboratory Model (ULM) – that can effectively model multidimensional laboratory data with missing values using the attention mechanism of the Set Transformer [23] and internal encodings of key parameters by the Large Language Model [24], [25].
- Applying this architecture to laboratory data, we built accurate models to predict system pathology that can result in low ferritin, high uric acid, cholesterol, and glucose levels based on CBC results and (if any) other blood tests. The model can be easily expanded for other tests (features).
- We made these models publicly available on <https://ulm.roslis.ru> to facilitate their use by Laboratory Information Systems throughout the world. The source

code developed in the work is also available on GitHub, <https://github.com/petilia2002/ferritin>.

II. METHODS

A. Dataset

Data used for this study were collected from drivers connecting the Laboratory Information System (RosLIS²) to analytical equipment in a number of commercial medical laboratories in Russia. These laboratories provide service for ordinary people, patients in hospitals, check-up inspections, making the data quite diverse and not biased to a population of narrow specialized organizations as is used in many contemporary studies; see Table I. The size of our database is 1 173 403 data rows and is 10 times larger compared to known studies.

As the data came from laboratories that analyze several thousands samples per day, technical errors (test samples, equipment, or quality control failures, not human samples (animals, dead organisms, etc.) are inevitable and a special procedure to eliminate clearly wrong values is needed as there are no standard intervals to exclude outliers taking into account that a severe pathology may have been observed. By developing this procedure, we also keep in mind an applicability domain estimation of the model to ensure that it can provide reliable results.

B. Applicability domain

Any mathematical model can only produce reliable results if the input data fall in the model’s applicability domain (AD), which usually includes ranges for all continuous parameters and predicted outcomes. Acceptable value ranges can be determined for each blood parameter by constructing a probability density function that captures variations in the parameter. The acceptable range is then set to include some predefined percentage of the area under the distribution curve. As our

²<https://roslis.ru>

database has more than one million rows and the data are supposed to be reliable, we set this threshold to a high value, 99.99%.

Our implementation of this algorithm includes building a histogram of all the values of a test, sorting it by them, finding the maximum, and traversing from this point in both directions until the fraction of the cumulative number of visited values is less than the threshold.

The ranges found are presented in Table II. Interestingly, the minimum acceptable value for HGB is 11 g/L as it is known that less than 10 g/L for the HGB test is not compatible with life. In contrast, the maximum value of NEUT is 100% seems to be unreliable since NEUT is a part of WBC and very rare a part could be the whole. We decided not to exclude these values because, from a numerical perspective, they would not induce any distortion.

In addition to ranges, there is also another condition that must be satisfied for a sample to be in the AD - the sum of all WBC parts from a 5-DIFF analyzer, namely BASO, EOS, LYMPH, MONO, and NEUT (or MID, GRA, and LYMP for 3-DIFF) must equal $100 \pm 4\%$ ³.

C. Universal Laboratory Model

Building a model on tabular data lacks information on the data itself. All the model knows is that there are some columns and that there are some values in these columns. But it has no clue what exactly each column means and how these columns are related or what objectives of the model are and how they possibly can relate between each other. In general, there is no context that the model can benefit from when training for a particular task. Without this context, deep learning models stagnate and are usually not better than classical machine learning algorithms like RF or SVM.

To catch up on these relations between input and output and provide context, an attention mechanism can be used. Similar column embeddings were introduced in [26] as trainable parameters that can improve the performance of the model. Instead, we use GPT-like embeddings of features that are fixed and propose to use the same-type embeddings for objectives of the model as well. Our approach uses the Encoder-Decoder pattern, Fig. 1. The attention block has three parameters, namely queries Q , keys K , and values V [24]. If X denotes input data and P - predicting features, then for the encoder $E(X)$, we use the so-called Self-Attention, where all the parameters are the same and correspond to the input of the model; for the decoder $D(P, X)$ the $Q = P$ and comes from the features the model is training to predict:

$$E(X) = \text{Attention}(Q = X, K = X, V = X) \quad (1)$$

$$D(P, X) = \text{Attention}(Q = P, K = I, V = I) \quad (2)$$

There is no natural order in placing test results on a sheet of a laboratory report except traditions and perhaps state regulations. If one casts away missing values, all position

³We can accept a small difference from 100% as sometimes the values are edited by hand.

Table II: Ranges of values of parameters, alphabetical order

#	Feature (Test)	Min	Max	Units	LOINC ^a
1	Age	1	110	years	—
2	Alanine transaminase, (ALT)	0.10	1721.40	U/L	1742-6
3	Albumin, (ALB)	0.04	60.14	g/L	1751-7
4	Aspartate aminotransferase, (AST)	0.30	1518.50	U/L	1920-8
5	Basophils, (BASO)	0.02	27.20	%	704-7
6	Cholesterol, (CHOL)	0.02	17.20	mmol/L	14647-2
7	Creatinine, (CREA)	0.30	1618.80	$\mu\text{mol/L}$	14682-9
8	C-reactive protein, (CRP)	0.01	250.95	mg/L	1988-5
9	Direct bilirubin, (BC)	0.04	238.40	$\mu\text{mol/L}$	29760-6
10	Eosinophils, (EOS)	0.09	43.80	%	711-2
11	Ferritin, (FER)	0.01	1667.20	$\mu\text{g/L}$	20567-4
12	Folic acid, (FOL)	0.56	330.11	ng/mL	2284-8
13	Gender, ($\{M = 1, F = 0\}$)	0	1	—	—
14	Glucose, (GLU)	0.01	26.77	mmol/L	14771-0
15	Granulocytes, (GRA)	14.7	94.7	%	19023-1
16	Hemoglobin, (HGB)	11.00	215.00	g/L	30350-3
17	Indirect bilirubin, (BU)	0.05	220.79	$\mu\text{mol/L}$	14630-8
18	Lactate dehydrogenase, (LDH)	2.00	4983.00	U/L	2532-0
19	Lymphocytes, (LYMPH)	0.10	90.70	%	737-7
20	Mean corpuscular volume, (MCV)	0.70	134.00	fL	71829-6
21	Middle-size Cells, (MID)	1.2	29.4	%	32155-4
22	Monocytes, (MONO)	0.10	55.40	%	5905-5
23	Neutrophils, (NEUT)	2.90	100.00	%	768-2
24	Platelets, (PLT)	1.00	1053.00	$10^9/\text{L}$	777-3
25	Red blood cells, (RBC)	0.24	8.21	$10^{12}/\text{L}$	789-8
26	Total bilirubin, (TBIL)	0.03	432.43	$\mu\text{mol/L}$	54363-7
27	Total protein, (PRO)	19.20	132.10	g/L	13980-8
28	Urea, (UREA)	0.50	67.50	mmol/L	22664-7
29	Uric acid, (URIC)	0.00	1.22	mmol/L	14933-6
30	Vitamin B12, (VB)	1.00	39833.00	pg/mL	2132-9
31	White blood cells, (WBC)	0.10	68.70	$10^9/\text{L}$	804-5

^aLogical Observation Identifiers Names and Codes – <https://loinc.org>

information is gone, and all is rest is a number of unordered sets. To convert a set to a vector suitable for machine learning, one has to make sure that an order of elements in this vector does not affect model's output.

Pooling techniques such as dimension-wise average or maximum of feature vectors are the usual choice to make this output invariant. The attention block (multi-head attention as well) [24] is equivariant, meaning that any permutation in the input gives the same permutation in the output. To make it invariant, it is possible to utilize a multi-head attention pooling [23] induced by queries Q , which are not dependent on the order of input and, therefore, could be learnable. By construction, ULM can work with sets of different sizes and

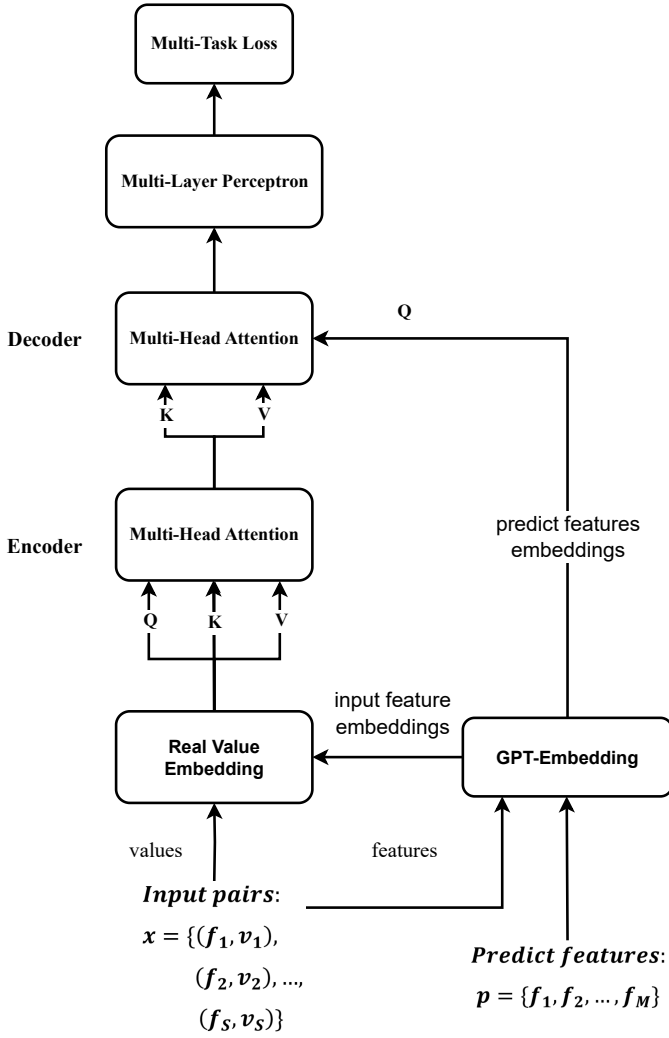


Fig. 1: ULM architecture: Encoder-Decoder pattern, Decoder uses predicting feature embeddings as queries (Q)

orders and produce a consistent prognosis.

For simplicity, we consider only one laboratory report (for one patient). Let $x = (f, v)$ denote a pair of feature values (test result), where f is a categorical feature of a set of all available features (tests), Table II (column #), $F = \{f_1, f_2, \dots, f_L\}$, including age, gender, and laboratory tests names⁴, L - number of all possible features (here 31); $v \in \mathbb{R}$ - is a numerical value of that feature. A subset $P \subset F$ contains the tests which abnormal values the model is going to predict. As we restrict ourselves to model only for tests (GLU, CHOL, FER, URIC), the cardinality of P is 4. A laboratory report consists of O results and could be expressed as $\mathbf{r} = \{x_1, x_2, \dots, x_O\}$, where $O \leq L$. Among these O pairs, we treat S pairs as training input data $\mathbf{x} = \{x_1, x_2, \dots, x_S \mid f \notin P\}$ and $S \leq L - |P|$ and the rest of \mathbf{r} is training output data $\mathbf{y} = \{\phi_1(x_1), \phi_2(x_2), \dots, \phi_M(x_M) \mid f \in P\}$, where

⁴Strictly speaking, age and gender have a different nature compared to laboratory tests though they could be examined and, therefore, treated as results as well.

Table III: Example records from our database, $N = 1\,173\,403$. Data to the left of the line is used as input, to the right – output. Empty cells depict missing values.

#	Gender	Age	HGB	RBC	MCV	...	UREA	CREA	GLU	CHOL	FER	URIC
1	F	55	128	4.48	85	.	3.4	75.2	5.37		8.67	0.247
2	F	33	137	4.22	97.2	.			4.4	5.44	124.1	0.521
3	M	66	146	4.29	100	.	109		8.24	5.29	94.6	0.314
.
N	M	31	141	4.65	86	.	3.2	86.7		3.98		0.327

$M = |P|$ and ϕ_i is a function that converts a real-value of a test i into a binary class as in Table IV.

Each row in Table III corresponds to \mathbf{r} ; \mathbf{x} is a part of \mathbf{r} placed to the left from the vertical line, and the source for \mathbf{y} is a part of \mathbf{r} placed to the right from this line. For sample #1 in Table III:

$$\begin{aligned} \mathbf{x}_1 &= \{(13, 0), (1, 55), (16, 128), (25, 4.48) \\ &\quad (20, 85), (28, 3.4), (7, 75.2)\} \\ \mathbf{y}_1 &= \{(14, 0), (11, 1), (29, 0)\} \end{aligned}$$

For different samples (patients), the O may vary as the doctors order different sets of tests depending on the patient's state, possible disease, or general objective of an inspection. Moreover, the training set is constructed so $\mathbf{y} \notin \emptyset$, while during inference \mathbf{y} must be empty or at least not equal to P .

We believe GPT-embeddings carry important information for the model and all rotations and squeezing of these embeddings destruct GPT knowledge, therefore, any matrix multiplication must be refrained. Only scalar multiplication and translation are left, so the output of Real Value Embeddings layer will be:

$$y_i = GPT_i * v_i + B \quad (3)$$

where GPT stands for GPT-embeddings of feature i , v_i is the value of that feature, B is a trainable translation parameter, and $*$ is scalar multiplication.

D. Data preparation and scaling

To form the dataset for training and testing, we apply the ranges in Table II to all available data and select only those results that fall in the ranges and contain values of at least one of the tests from P , that is glucose, cholesterol, ferritin, or uric acid. For every value v in \mathbf{x} we apply the logarithmic transformation since there are no zero values and then scale them to the interval $[0.1, 0.9]$ for normalization.

E. Training and validation protocols

All models were trained with the same protocol. The dataset was randomly divided 80% for training and 20% for testing. Training was controlled by the Adam optimization algorithm with early stopping and automatic reduction of the learning rate techniques. The total number of epochs was limited to 1 000, and the statistics in Table IV are the result of three

different training runs for each variant. For standard Multilayer Perceptron (MLP) models, the optimization started with the learning rate 10^{-3} , for ULM 10^{-4} . The batch size was 32; to support batch mode, we used the usual Transformer Masks for the encoder and decoder. Tests embeddings were calculated using the open Yandex GPT service [25]. All code is written in Python3 programming language. The Keras framework [27] for machine learning with TensorFlow [28] backend is used.

III. RESULTS

In our first experiments, we used one output (refer to MLP_B in Table IV) to build classification models that are simple 3-layer neural networks with architecture 27 x 256 (relu) x 1 (sigmoid) and binary cross-entropy loss. We did not observe any dependency on the number of hidden neurons or activation function in the middle layer. Then we switched to multitask [29] settings (MLP_M) when the output comprises several masked values with masked categorical cross-entropy loss and the architecture is 27 x 256 (relu) x 4 (sigmoid); but no notable improvements were found. For all MLP models, we varied training settings, activation functions, normalization schemes, applied different residual connections, built numerical embeddings [30], performed internal ensemble [31], but all our efforts did not affect the performance of the models; small fluctuations in AUC - no better, no worse.

The next step was to use ULM for the same task. ULM uses a multi-head attention encoder and decoder with a number of heads equal to 8 and key dimension 16. The dropout is set to 0.1. ULM has the same output and loss as MLP_M. ULM improved the statistics for all modeled tests, Table IV. The ROC curves for the ULM model are shown in Fig. 2, confusion matrix presented in Table V.

The improvement for Ferritin is the smallest 0.8% (but significant), as this test has a strong relation to only CBC (especially HGB) and, as we see, additional biochemical data does not substantially change the situation. The AUC for Glucose and Cholesterol outputs got $\sim 2\%$ gain, while the increase for Uric acid jumps up to 8%. One possible explanation is that high values of Uric acid in the blood cause gout, and during attack a patient has severe pain in the joints that affects all functions of the organisms and is reflected in the change in CBC and other biochemical tests.

Table IV: Models' AUC values (%) (the bigger the better)

Test	Threshold ^a	Size	MLP_B ^b	MLP_M ^c	ULM
GLU	≥ 7 mmol/L	680 502	80.2 \pm 0.1	80.4 \pm 0.1	82.4\pm0.1
CHOL	≥ 5.2 mmol/L	248 806	75.7 \pm 0.1	75.6 \pm 0.1	77.9\pm0.1
FER	≤ 12 ng/mL	105 928	88.9 \pm 0.2	88.7 \pm 0.1	89.7\pm0.1
URIC	M: ≥ 0.48 mmol/L F: ≥ 0.38 mmol/L	92 901	71.5 \pm 0.2	71.5 \pm 0.2	79.7\pm0.2

^aIf the condition is true, the sample is treated as positive. ^bBinary classification: the model has one output and each analyte is modeled independently of the others. ^cMulticlass classification: the model has 4 outputs, e.g. all analytes are modeled simultaneously.

Table V: Parameters of the ULM model

Parameter	Glucose	Cholesterol	Ferritin	Uric acid
True Positive (TP)	5 337	22 616	2 066	1 940
True Negative (TN)	101 194	16 276	16 488	12 100
False Positive (FP)	34 454	6 928	3 850	4 965
False Negative (FN)	1 943	9 657	520	901
Accuracy (%)	74.5	70.1	80.9	70.5
Sensitivity (%)	73.3	70.1	79.9	68.3
Specificity (%)	74.6	70.1	81.1	70.9

IV. CONCLUSIONS AND OUTLOOK

In this article, we propose the Universal Laboratory Model that uses context GPT-like embeddings for both input and output and can effectively work with different size sets thanks to the invariant possibility in the encoder-decoder architecture. In four laboratory cases, we demonstrate that the use of ULM can improve model performance compared to standard MLP, and the models can be used in laboratory diagnostics.

Still there are open questions. What is a right threshold to mark a positive case? As shown in Table I, different groups define different thresholds leading to inconsistent statistics. Moreover, every reading has its own error and is influenced by more broad biological variation. We think incorporation of this additional data could also be beneficial for the models making them more consistent.

Another problem is data. As we stress, many groups use a limited number of data points to build and compare models. Usually the origin of these data is one hospital and therefore is subjected to limited groups of patient that already have some kind of health disorders, so the data contains more positive




Fig. 2: ROC-curves for ULM model. Positive class means prediction of abnormal value.

examples than negative and the model statistics are biased to positive cases showing remarkably good values. Utilizing datasets from different hospitals and even countries raise more complex tasks: How to deal with state regulations that personal data (including any medical records) originating within a country must stay within this country? Techniques such as federated learning [32] might help, but despite the success in theory and even practice in decentralized computation and data storage, its use is challenging and limited.

From the patient's perspective, the question "what to do with a prognosis in practice" arises. Simple commenting on the result on a sheet of laboratory report like "You may have high uric acid" is useless, as it brings more questions and concerns to the patient than helps her, especially if the prognosis is wrong. Most physicians would not recommend such comments. Another possible solution that we propose is to inform the physician through any channels of integration between medical information systems (MIS). In this scenario, when a doctor views a laboratory report, MIS sends data to the AI service and presents its recommendation, but the final action depends on the doctor.

Finally, we suggest that laboratories use these models to automatically add new tests that are predicted as positive outcomes. This would work as follows; when a laboratory system receives results, it asks for the AI service, and if there are positive predictions and biological material from the same patient is also available, the corresponding tests are added and performed automatically. In this scenario, the patient receives a laboratory report with additional test results that were not prescribed, and the doctor can start appropriate treatment the very next time she sees the patient.

REFERENCES

- [1] World Health Organization, "The selection and use of essential in vitro diagnostics: Report of the third meeting of the strategic advisory group of experts on in vitro diagnostics," World Health Organization, Tech. Rep. 1031, 2021. [Online]. Available: <https://iris.who.int/bitstream/handle/10665/339064/9789240019102-eng.pdf?sequence=1>
- [2] M. A. Santos-Silva, N. Sousa, and J. C. Sousa, "Artificial intelligence in routine blood tests," *Frontiers in Medical Engineering*, vol. 2, 2024. [Online]. Available: <https://www.frontiersin.org/journals/medical-engineering/articles/10.3389/fmede.2024.1369265>
- [3] I.-H. Seo and Y. Lee, "Usefulness of complete blood count (cbc) to assess cardiovascular and metabolic diseases in clinical settings: A comprehensive literature review," *Biomedicine*, vol. 10, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:253148030>
- [4] S. Kurstjens, T. de Bel, A. van der Horst, R. Kusters, J. Krabbe, and J. van Balveren, "Automated prediction of low ferritin concentrations using a machine learning algorithm," *Clinical Chemistry and Laboratory Medicine (CCLM)*, vol. 60, no. 12, pp. 1921–1928, 2022. [Online]. Available: <https://doi.org/10.1515/cclm-2021-1194>
- [5] Y. Luo, P. Szolovits, A. S. Dighe, and J. M. Baron, "Using machine learning to predict laboratory test results," *American Journal of Clinical Pathology*, vol. 145, no. 6, pp. 778–788, 06 2016. [Online]. Available: <https://doi.org/10.1093/ajcp/aqw064>
- [6] S. Pullakhandam and S. McRoy, "Classification and Explanation of Iron Deficiency Anemia from Complete Blood Count Data Using Machine Learning," *BioMedInformatics*, vol. 4, no. 1, pp. 661–672, 2024. [Online]. Available: <https://www.mdpi.com/2673-7426/4/1/36>
- [7] G. Cardozo, G. B. Pintarelli, G. R. Andreis, A. C. W. Lopes, and J. L. B. Marques, "Use of machine learning and routine laboratory tests for diabetes mellitus screening," *BioMed Research International*, vol. 2022, no. 1, p. 8114049, 2022. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1155/2022/8114049>
- [8] S. L. Cichosz, C. Bender, and O. Hejlesen, "A comparative analysis of machine learning models for the detection of undiagnosed diabetes patients," *Diabetology*, vol. 5, no. 1, pp. 1–11, 2024. [Online]. Available: <https://www.mdpi.com/2673-4540/5/1/1>
- [9] Y. Zhang, L. Zhang, H. Lv, and G. Zhang, "Ensemble machine learning prediction of hyperuricemia based on a prospective health checkup population," *Frontiers in Physiology*, vol. 15, 2024. [Online]. Available: <https://www.frontiersin.org/journals/physiology/articles/10.3389/fphys.2024.1357404>
- [10] S. Lee, E. K. Choe, and B. Park, "Exploration of machine learning for hyperuricemia prediction models based on basic health checkup tests," *Journal of Clinical Medicine*, vol. 8, no. 2, 2019. [Online]. Available: <https://www.mdpi.com/2077-0383/8/2/172>
- [11] Y. Gao, S. Jia, D. Li, C. Huang, Z. Meng, Y. Wang, M. Yu, T. Xu, M. Liu, J. Sun, Q. Jia, Q. Zhang, Y. Gao, K. Song, X. Wang, and Y. Fan, "Prediction model of random forest for the risk of hyperuricemia in a chinese basic health checkup test," *Bioscience Reports*, vol. 41, no. 4, p. BSR20203859, 04 2021. [Online]. Available: <https://doi.org/10.1042/BSR20203859>
- [12] H. Zhou, F. Liu, B. Gu, X. Zou, J. Huang, J. Wu, Y. Li, S. S. Chen, P. Zhou, J. Liu, Y. Hua, C. Mao, C. You, X. Wu, Y. Zheng, L. Clifton, Z. Li, J. Luo, and D. A. Clifton, "A survey of large language models in medicine: Progress, application, and challenge," 2024. [Online]. Available: <https://arxiv.org/abs/2311.05112>
- [13] Y. Gorishniy, I. Rubachev, V. Khurlov, and A. Babenko, "Revisiting deep learning models for tabular data," 2023. [Online]. Available: <https://arxiv.org/abs/2106.11959>
- [14] M. B. Sampa, M. N. Hossain, M. R. Hoque, R. Islam, F. Yokota, M. Nishikitani, and A. Ahmed, "Blood uric acid prediction with machine learning: Model development and performance comparison," *JMIR Med Inform*, vol. 8, no. 10, p. e18331, Oct 2020. [Online]. Available: <https://medinform.jmir.org/2020/10/e18331>
- [15] R. Vohra, A. Hussain, A. K. Dudyala, J. Pahareeya, and W. Khan, "Multi-class classification algorithms for the diagnosis of anemia in an outpatient clinical setting," *PLOS ONE*, vol. 17, no. 7, pp. 1–18, 07 2022. [Online]. Available: <https://doi.org/10.1371/journal.pone.0269685>
- [16] T. Emmanuel, T. Maupong, D. Mpoeleng, T. Semong, B. Mphago, and O. Tabona, "A survey on missing data in machine learning," *Journal of Big data*, vol. 8, pp. 1–37, 2021.
- [17] L. O. Joel, W. Doorsamy, and B. S. Paul, "A review of missing data handling techniques for machine learning," *International Journal of Innovative Technology and Interdisciplinary Sciences*, vol. 5, no. 3, pp. 971–1005, 2022.
- [18] S. Katoch, S. S. Chauhan, and V. Kumar, "A review on genetic algorithm: past, present, and future," *Multimedia tools and applications*, vol. 80, pp. 8091–8126, 2021.
- [19] A. Sulejmani and O. Koça, "Development of optimal transmission rate of the kinematic chain by using genetic algorithms coded in mathcad," *International Journal of Innovative Technology and Interdisciplinary Sciences*, vol. 4, no. 4, pp. 792–803, 2021.
- [20] S. Yoon and S. Sull, "Gamin: Generative adversarial multiple imputation network for highly missing data," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 8456–8464.
- [21] J. Kim, D. Tae, and J. Seok, "A survey of missing data imputation using generative adversarial networks," in *2020 International conference on artificial intelligence in information and communication (ICAIIIC)*. IEEE, 2020, pp. 454–456.
- [22] E. El-Kenawy and E. SM, "A machine learning model for hemoglobin estimation and anemia classification," *International Journal of Computer Science and Information Security (IJCSIS)*, vol. 17, no. 2, pp. 100–108, 2019.
- [23] J. Lee, Y. Lee, J. Kim, A. R. Kosiorek, S. Choi, and Y. W. Teh, "Set transformer: A framework for attention-based permutation-invariant neural networks," 2019. [Online]. Available: <https://arxiv.org/abs/1810.00825>
- [24] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is All you Need," in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017.
- [25] [Online]. Available: <https://yandex.cloud/en-ru/services/yandexgpt>

- [26] X. Huang, A. Khetan, M. Cvitkovic, and Z. Karnin, "Tabtransformer: Tabular data modeling using contextual embeddings," 2020. [Online]. Available: <https://arxiv.org/abs/2012.06678>
- [27] F. Chollet *et al.*, "Keras," <https://keras.io>, 2015.
- [28] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015, software available from tensorflow.org. [Online]. Available: <https://www.tensorflow.org/>
- [29] S. Sosnin, M. Vashurina, M. Withnall, P. Karpov, M. Fedorov, and I. V. Tetko, "A survey of multi-task learning methods in chemoinformatics," *Molecular Informatics*, vol. 38, no. 4, p. 1800108, 2019. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/minf.201800108>
- [30] Y. Gorishniy, I. Rubachev, and A. Babenko, "On embeddings for numerical features in tabular deep learning," 2023. [Online]. Available: <https://arxiv.org/abs/2203.05556>
- [31] Y. Gorishniy, A. Kotelnikov, and A. Babenko, "Tabm: Advancing tabular deep learning with parameter-efficient ensembling," 2025. [Online]. Available: <https://arxiv.org/abs/2410.24210>
- [32] B. Yurdem, M. Kuzlu, M. K. Gullu, F. O. Catak, and M. Tabassum, "Federated learning: Overview, strategies, applications, tools and future directions," *Heliyon*, vol. 10, no. 19, p. e38137, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2405844024141680>