FREE-TEXT AND STRUCTURED CLINICAL TIME SERIES FOR PATIENT OUTCOME PREDICTIONS

Emilia Apostolova PhD[1], Joe Morales[2], Ioannis Koutroulis, MD PhD[3] and Tom Velez, PhD, JD[2]

[1]Language.ai, Chicago, IL, USA
[2]Computer Technology Associates, Ridgecrest, CA, USA
[3]Children's National Health System, Washington, DC, USA

ABSTRACT

While there has been considerable progress in building deep learning models based on clinical time series data, overall machine learning (ML) performance remains modest. Typical ML applications struggle to combine various heterogenous sources of Electronic Medical Record (EMR) data, often recorded as a combination of free-text clinical notes and structured EMR data. The goal of this work is to develop an approach for combining such heterogenous EMR sources for time-series based patient outcome predictions. We developed a deep learning framework capable of representing free-text clinical notes in a low dimensional vector space, semantically representing the overall patient medical condition. The free-text based time-series vectors were combined with time-series of vital signs and lab results and used to predict patients at risk of developing a complex and deadly condition: acute respiratory distress syndrome. Results utilizing early data show significant performance improvement and validate the utility of the approach.

KEYWORDS

Natural Language Processing; Clinical NLP; Time-series data; Machine Learning; Deep Learning; Free-text and structured data; Clinical Decision Support; ARDS; COVID-19

1. INTRODUCTION

Deep Learning utilizing Electronic Medical Record (EMR) data for medical diagnosis/outcome predictions is an active and promising field of research. The interest in the topic has been spurred by the combination of a number of contributing factors. On one hand, there is the availability and sheer abundance of EMR data: in the US alone, the Centers for Disease Control and Prevention report more than 800 million physician office visits annually, most associated with digital EMR data [1]. This, combined with the practical significance of medical AI, advances in deep learning, and the availability of powerful and inexpensive computing resources, has led to an abundance of clinical prediction models derived to predict various medical outcomes with limited clinical success [2,3].

More recently, the utility of time-series EMR data has been explored for improved deep learning predictions, as traditional ML on the entire time series is often infeasible as each data point would be handled as a separate feature introducing dimensionality problems. Patient visit EMR data, such as vital signs, lab results, clinical notes, etc., is typically time-stamped, and, intuitively,

David C. Wyld et al. (Eds): ITCSE, NLCA, ICAIT, CAIML, ICDIPV, CRYPIS, WiMo - 2020
pp. 129-136, 2020. CS & IT - CSCP 2020 DOI: 10.5121/csit.2020.100511

human experts often base their judgments on the temporal relations of various variables. For example, a time series showing non-improving respiratory rate values, despite medical interventions, may serve as a sign of clinical deterioration, as well as other outcomes (e.g. mortality risk). Recently, time-series-based ML effort attempts to learn such temporal clinical knowledge and multi-task inference using a clinical time series benchmark dataset [4,5], derived from the publicly available Medical Information Mart for Intensive Care (MIMIC-III) database [6]. The dataset contains time series data for 17 selected clinical variables containing more than 31 million clinical events. The benchmark tasks consist of in-hospital mortality prediction, decompensation prediction, length-of-stay prediction, and phenotype classification. Harutyunyun et al. [4] also built several baseline ML models on the benchmark, including several LSTM-based models.

The clinical time series benchmark was also used on the task of medical diagnosis code prediction. Inspired by the success of embeddings combined with recurrent networks in NLP, Lipton et al. [7] built an LSTM-based diagnosis prediction model utilizing 13 time series variables used to predict 128 common diagnosis codes. More recently, [8] developed an attention model outperforming the LSTM models on a number of the MIMIC time series benchmark tasks.
The performance of the proposed time-series-based ML models, however, is quite limited. For example, the best achieved F1-scores on the multi-label diagnosis task described by Harutyunyun et al. [4] are 0.29 and 0.15, micro and macro F1-scores respectively. Human expert diagnosis coding significantly outperforms the proposed models simply because clinicians have access to additional patient data (outside the time series of 13 vital signs and lab results variables), that provides rich patient medical context. In particular, clinicians have access to the clinical free-text notes, that include information such as the patient medical history, family history, the reason for the visit, signs, symptoms, findings, social history, etc.

2. TASK DEFINITION

The goal of this study is to better approximate the clinical information used by human experts, by combining time series structured data (lab results and vital signs), with time series free text data (clinical notes). In particular, we focus on the early identification (within 24 hours of admission) of patients at risk of developing ARDS (Acute Respiratory Distress Syndrome). The condition is characterized by the development of acute dyspnea at rest, hypoxemia, and alveolar infiltrates on chest imaging within hours to days of an inciting event such as viral pneumonia. ARDS is a significant cause of morbidity and mortality in the USA and worldwide [9,10] and is the principal cause of COVID-19 associated mortality.
In a reported Wuhan case series, among the 201 patients with confirmed COVID-10, 41.8% developed ARDS and among these patients 52.4% died. It has been reported that in general ARDS caused by COVID-19 results in 2.3% mortality rate of diagnosed cases [11]. Early recognition can limit the propagation of lung injury and significantly improve patient outcomes [12]. Similar to other acute conditions, predicting ARDS is a difficult task even for human experts, as the condition is often confounded by cardiogenic factors, and, at the same time, is highly heterogeneous [13] and COVID-induced ARDS is atypical. ARDS involves the interaction of multiple risk factors, past history, and current conditions, signs, and symptoms, and thus structured time series data, without access to the free-text patient context will be insufficient in judging ARDS outcomes.

3. RELATED WORK

The literature related to this study falls into two categories: combining free-text and structured EMR data for clinical outcome predictions, and machine learning with clinical time series data.

A large volume of literature on combining structured and free-text EMR data pre-processes the free-text data by applying some information extraction (IE) technique, typically medical concept detection [14,15,16]. The majority of approaches extract UMLS or SNOMED-CT concepts from free-text with their negation status with various off-the-shelf tools [15, 17, 18, 19, 20].

More recently, Miotto et al. [21] built Deep Patient representations utilizing structured EMR variables and notes converted to a set of concepts using traditional methods. Shickel et al. [3] present a survey of various deep learning techniques, the majority of which focus on structured EMR data. In addition, a number of deep learning studies explore pre-training on diagnosis and procedure code embeddings [22,23,24].

In terms of utilizing time series data, deep learning techniques have been explored extensively, typically focusing only on structured EMR data. In addition to the studies focusing on LSTM and transformer architectures for clinical time series described in Section Introduction [4,7,8], a number of studies explore clinical time series for patient outcome prediction. For example, Choi et al. [25] develop a temporal model using recurrent neural networks (RNN) and time-stamped structured EMR data. Similarly, Lipton et al. [26] explore the modeling of missing time series data with RNNs. Choi, et al. [27] developed a reverse time attention model, so that recent clinical visits are likely to receive higher attention. Razavian et al. [28] built multi-task disease onsets prediction utilizing LSTM and CNN on common lab test results. Nguyen et al. [29] built a CNN model using coded EMR data, combined with coded time separators, such as [1-3 months], [6-12 months], etc. Xu et al. [30] developed a recurrent multi-channel attention model combining various clinical sources of time-series records including waveform and numeric data.

The main contribution of this work is the low-dimensional vector space representation of free-text, that can be combined with structured EMR data in the context of time-series based clinical outcome prediction.

4. METHOD

4.1. Dataset

Clinical encounter data of adult patients was extracted from the MIMIC3 Intensive Care Unit (ICU) database [6]. MIMIC3 consists of retrospective ICU encounter data of patients admitted into Beth Israel Deaconess Medical Center from 2001 to 2012. MIMIC3 includes time series data recorded in the EMR during encounters (e.g. vital signs/diagnostic laboratory results, free text clinical notes, medications, procedures, etc.). The dataset contains data associated with over 58,000 ICU visits, including over 2 million free-text clinical notes.

For this study, in accordance with previous literature [31], we identified ARDS for adult patients older than 18 years with ICD-9 codes for severe acute respiratory failure and use of continuous invasive mechanical ventilation, excluding those with codes for acute asthma, COPD and CHF exacerbations. This resulted in 4,624 ARDS cases from a total of 48,399 adult ICU admissions. The ICU mortality rate in this population was approximately 59%.

4.2. Structured Data Time Series

Time series data was collected over the first 24 hours of ICU admission. The first 24-hour timeframe was chosen, as it has been reported that ARDS develops at a median of 30 hours after hospital admission [32]. Thus, a 24-hour window provides for the gathering of enough data, while at the same time is early enough for real-time clinical decision support (CDS). Time series were created in 4-hour windows (averaging values if more than one present).

Utilizing clinical knowledge and ML experiments on the dataset, we identified the most informative vital and lab result variables in the context of ARDS outcome predictions. They include 'Bicarbonate','Systolic blood pressure (noninvasive)', 'Tidal Volume (set)', 'Partial pressure of carbon dioxide (arterial)', 'Monocytes', 'Partial pressure of oxygen', 'Lactate dehydrogenase', 'Urine output', 'Calcium (total)', 'Mean corpuscular hemoglobin concentration', 'Lymphocytes', 'Respiratory rate', 'Glascow coma scale verbal response', 'Minute Volume', 'Phosphate', 'Respiratory rate (total)', 'Heart Rate', 'Mean Airway Pressure', 'PEEP set', 'Diastolic blood pressure (noninvasive)', as well as 'age' and 'bmi'.

4.3. Free-text Data Time Series

In addition to structured data, we also included time-series of time-stamped clinical notes. In particular, we focused on nursing notes, as they are available early (within 24 hours), unlike, for example, discharge notes available at the end of the patient stay. Nursing notes also contain a comprehensive summary of the patient history and present condition. Similarly to the structured data time series, free-text time series were created in 4-hour windows. Ideally, we would like the free-text notes to be converted to a low dimensional vector space, semantically representing the overall patient medical condition, including medical history and present illness and symptoms.
It has been noted that clinicians viewing properly coded patient diagnosis codes (ICD9 and ICD10 codes) are typically capable of deducing the overall condition, history, and risk factors associated with a patient [33]. Diagnosis codes are used to describe information, such as current diagnoses, signs and symptoms, history and chronic conditions, past and current treatments / procedures, age group and/or susceptibilities, expected outcome, patient social history, the reason for the visit, etc. Intuitively, the totality of patient’s diagnosis codes represent a meaningful medical summary of the patient. However, real-time CDS systems, such as predicting ARDS outcome within 24 hours of admission, won't have access to the full set of the patient's ICD codes, which are typically entered at a later time.

At the same time, it has been suggested that the medical code co-occurrence of diagnosis can be exploited to generate low-dimensional representations of ICD codes [22-24] that may facilitate EMR data-based exploratory analysis and predictive modeling [33-35]. Building upon this work, we built a deep learning model trained to predict the patient’s ICD code embeddings from nursing notes and thus create a low dimensional vector space semantically representing the overall patient medical condition: *Patient Context Vectors* (PCV).

The full set of MIMIC3 nursing notes (1,081,176 free-text notes) were used as a pre-training step in a model trained to predict the patient's averaged ICD-code embeddings: PCVs. The optimum size of the ICD-code embedding vectors was determined to be 50. Experiments with two deep learning networks were performed. In both cases, the architectures utilized were similar to typical deep learning text classification networks, with the difference that the target prediction is not probabilities on a set of categories (soft-max loss function), but an ICD-code embedding vector (multi target regression with mean squared error loss function). In both cases, the input texts were truncated/padded to a length 400 tokens, the last linear layer of size 50 used loss function of mean squared error, the Adam optimizer was used with batch size of 32, trained on 3 epochs. A word-level CNN model [36], consisting of a convolutional, max Pooling layers, followed by 2 hidden layers of size 500 achieved a mean squared error of 0.18 on the test set. A fine-tuned Bert base model [37] achieved a mean squared error of 0.13 on the test set. Both models were used to convert test notes into to a low dimensional vector space (PCVs of size 50), semantically representing the overall patient medical condition.

4.4. Time-series ARDS Predictions

Time series data was collected from 6 four-hour windows following ICU admission. Each time series step contains values from 20 structured variables and nursing notes represented as PCVs of size 50, i.e. each time series step contains a total of 70 variables, with a total of 6 time series per visit. The representation is analogous to text classification representations, with word embeddings of size 70 and text length of size 6. All structured variables were normalized, and missing values were replaced with an indicator variable. A basic LSTM network was trained with an LSTM layer of size 200, followed by a dense layer with binary cross-entropy loss. The network used the Adam optimizer, LSTM-layer 0.3 dropout and L2 regularizer. Instances were weighted to accommodate for the unbalanced dataset (4,624 ARDS cases from a total of 48,399 adult ICU admissions). 10-fold cross-validation results are shown in Table 1.

Table 1. 10-fold cross-validation results for predicting ARDS outcomes from 6 time series steps (within 24 hours of admission). Structured: Structured data representing 20 vital signs and lab results; Structured + CNN PCV: Structured data and Patient Context Vectors of size 50 pre-trained using word-level CNN on all MIMIC3 nursing notes; Structured + Bert PCV: Structured data and Patient Context Vectors of size 50 pre-trained using Bert base model fine-tuning on all MIMIC3 nursing notes.

Time Series Data	Precision	Recall	F1-score
Structured	30.6	64	41.4
Structured + CNN PCV	36.3	65	46.6
Structured + Bert PCV	38.6	66	48.7

Results suggest that including free-text time-series data significantly outperforms predictions based exclusively on structured lab and vital signs. The addition of the CNN and the Bert-based models outperformed the baseline LSTM using only structured data by 5.2 and 7.3 F1-score absolute percent points respectively. Not surprisingly, the transformer-based Bert pre-trained model outperformed the word-level CNN results (by 2.1 F1-score absolute percent points). It is not clear how the results compare to human expert performance, as the outcome variable is based on subsequent ARDS outcome, and not on human judgements made within 24 hours of admission. It is likely that human expert ARDS predictions utilizing early admission data might also exhibit relatively low F-scores, as ARDS is an extremely challenging condition, requiring knowledge and inference based on complex geno- and pheno-type interactions.

5. CONCLUSIONS

This study focused on combining time-series lab results and vital signs EMR data, with free-text clinical notes time series attempting to capture patient medical context information. Our end goal was to predict early (within 24 hours) the development of an acute condition (ARDS), a task that is challenging even for clinical experts as it requires thorough knowledge and understanding of the patient's geno- and pheno-type, combined with the temporal monitoring of various tests and signs. Results suggest that the encoding and addition of the information present in free-text notes improved substantially the overall model performance.

REFERENCES

[1] P. Rui and T. Okeyode, “National ambulatory medical care survey: 2016 national summary tables,” 2016.

[2] A. Adibi, M. Sadatsafavi, and J. Ioannidis, “Validation and utility testing of clinical prediction models: Time to change the approach,” Jama, 2020. [Online]. Available: https://doi.org/10.1001/jama.2020.1230

[3] B. Shickel, P. J. Tighe, A. Bihorac, and P. Rashidi, “Deep ehr: a survey of recent advances in deep learning techniques for electronic health record (ehr) analysis,” IEEE journal of biomedical and health informatics, vol. 22, no. 5, pp. 1589–1604, 2018.

[4] H. Harutyunyan, H.Khachatrian, D.C.Kale, G.V.Steeg,and A.Galstyan,“Multitask learning and benchmarking with clinical time series data,” arXiv preprint arXiv:1703.07771, 2017.

[5] H. Harutyunyan, H. Khachatrian, D. C. Kale, G. Ver Steeg, and A. Galstyan, “Multitask learning and benchmark- ing with clinical time series data,” Scientific data, vol. 6, no. 1, pp. 1–18, 2019.

[6] A. E. Johnson, T. J. Pollard, L. Shen, H. L. Li-wei, M. Feng, M. Ghassemi, B. Moody, P. Szolovits, L. A. Celi, and R. G. Mark, “Mimic-iii, a freely accessible critical care database,” Scientific data, vol. 3, p. 160035, 2016.

[7] Z. C. Lipton, D. C. Kale, C. Elkan, and R. Wetzel, “Learning to diagnose with lstm recurrent neural networks,” arXiv preprint arXiv:1511.03677, 2015.

[8] H. Song, D. Rajan, J. J. Thiagarajan, and A. Spanias, “Attend and diagnose: Clinical time series analysis using attention models,” in Thirty-second AAAI conference on artificial intelligence, 2018.

[9] J. Ma´ca, O. Jor, M. Holub, P. Sklienka, F. Bursˇa, M. Burda, V. Janout, and P. Sˇevcˇ´ık, “Past and present ards mortality rates: a systematic review,” Respiratory care, vol. 62, no. 1, pp. 113–122, 2017.

[10] G. Bellani, J. G. Laffey, T. Pham, E. Fan, L. Brochard, A. Esteban, L. Gattinoni, F. Van Haren, A. Larsson, D. F. McAuley et al., “Epidemiology, patterns of care, and mortality for patients with acute respiratory distress syndrome in intensive care units in 50 countries,” Jama, vol. 315, no. 8, pp. 788–800, 2016.

[11] Z. Wu and J. M. McGoogan, “Characteristics of and important lessons from the coronavirus disease 2019 (covid- 19) outbreak in china: summary of a report of 72 314 cases from the chinese center for disease control and prevention,” Jama, 2020.

[12] E. Fan, L. Del Sorbo, E. C. Goligher, C. L. Hodgson, L. Munshi, A. J. Walkey, N. K. Adhikari, M. B. Am- ato, R. Branson, R. G. Brower et al., “An official american thoracic society/european society of intensive care medicine/society of critical care medicine clinical practice guideline: mechanical ventilation in adult patients with acute respiratory distress syndrome,” American journal of respiratory and critical care medicine, vol. 195, no. 9, pp. 1253–1263, 2017.

[13] C. M. Shaver and J. A. Bastarache, “Clinical and biological heterogeneity in acute respiratory distress syndrome: direct versus indirect lung injury,” Clinics in chest medicine, vol. 35, no. 4, pp. 639–653, 2014.

[14] S. DeLisle, B. South, J. A. Anthony, E. Kalp, A. Gundlapallli, F. C. Curriero, G. E. Glass, M. Samore, and T. M. Perl, “Combining free text and structured electronic medical record entries to detect acute respiratory infections,” PloS one, vol. 5, no. 10, p. e13377, 2010.

[15] H. Zheng, H. Gaff, G. Smith, and S. DeLisle, “Epidemic surveillance using an electronic medical record: an empiric approach to performance improvement,” PloS one, vol. 9, no. 7, p. e100845, 2014.

[16] E. Ford, J. A. Carroll, H. E. Smith, D. Scott, and J. A. Cassell, “Extracting information from the text of electronic medical records to improve case detection: a systematic review,” Journal of the American Medical Informatics Association, vol. 23, no. 5, pp. 1007–1015, 2016.

[17] A. V. Gundlapalli, B. R. South, S. Phansalkar, A. Y. Kinney, S. Shen, S. Delisle, T. Perl, and M. H. Samore, “Application of natural language processing to va electronic health records to identify phenotypic characteristics for clinical and research purposes,” Summit on translational bioinformatics, vol. 2008, p. 36, 2008.

[18] R.J.Carroll,A.E.Eyler,andJ.C.Denny,“Na¨ıve electronic health record phenotype identification for rheumatoid arthritis,” in AMIA annual symposium proceedings, vol. 2011. American Medical Informatics Association, 2011, p. 189.

[19] S.Karnik, S.L.Tan, B.Berg, I.Glurich, J.Zhang, H.J.Vidaillet, C.D.Page, and R.Chowdhary,"Predicting atrial fibrillation and flutter using electronic health records," in Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE. IEEE, 2012, pp. 5562–5565.

[20] A. N. Ananthakrishnan, T. Cai, G. Savova, S.-C. Cheng, P. Chen, R. G. Perez, V. S. Gainer, S. N. Murphy, P. Szolovits, Z. Xia et al., "Improving case definition of crohn's disease and ulcerative colitis in electronic medical records using natural language processing: a novel informatics approach," Inflammatory bowel diseases, vol. 19, no. 7, pp. 1411–1420, 2013.

[21] R. Miotto, L. Li, B. A. Kidd, and J. T. Dudley, "Deep patient: an unsupervised representation to predict the future of patients from the electronic health records," Scientific reports, vol. 6, p. 26094, 2016.

[22] Y. Choi, C. Y.-I. Chiu, and D. Sontag, "Learning low-dimensional representations of medical concepts," AMIA Summits on Translational Science Proceedings, vol. 2016, p. 41, 2016.

[23] E. Choi, M. T. Bahadori, E. Searles, C. Coffey, M. Thompson, J. Bost, J. Tejedor-Sojo, and J. Sun, "Multi-layer representation learning for medical concepts," in Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. ACM, 2016, pp. 1495–1504.

[24] D. Kartchner, T. Christensen, J. Humpherys, and S. Wade, "Code2vec: Embedding and clustering medical di- agnosis data," in Healthcare Informatics (ICHI), 2017 IEEE International Conference on. IEEE, 2017, pp. 386–390.

[25] E. Choi, M. T. Bahadori, A. Schuetz, W. F. Stewart, and J. Sun, "Doctor ai: Predicting clinical events via recurrent neural networks," in Machine Learning for Healthcare Conference, 2016, pp. 301–318.

[26] Z. C. Lipton, D. C. Kale, and R. Wetzel, "Modeling missing data in clinical time series with rnns," arXiv preprint arXiv:1606.04130, 2016.

[27] E. Choi, M. T. Bahadori, J. Sun, J. Kulas, A. Schuetz, and W. Stewart, "Retain: An interpretable predictive model for healthcare using reverse time attention mechanism," in Advances in Neural Information Processing Systems, 2016, pp. 3504–3512.

[28] N. Razavian, J. Marcus, and D. Sontag, "Multi-task prediction of disease onsets from longitudinal laboratory tests," in Machine Learning for Healthcare Conference, 2016, pp. 73–100.

[29] P. Nguyen, T. Tran, N. Wickramasinghe, and S. Venkatesh, "Deepr: A convolutional net for medical records," IEEE journal of biomedical and health informatics, vol. 21, no. 1, pp. 22–30, 2017.

[30] Y. Xu, S. Biswal, S. R. Deshpande, K. O. Maher, and J. Sun, "Raim: Recurrent attentive and intensive model of multimodal patient monitoring data," in Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2018, pp. 2565–2573.

[31] C. Bime, C. Poongkunran, M. Borgstrom, B. Natt, H. Desai, S. Parthasarathy, and J. G. Garcia, "Racial differ- ences in mortality from severe acute respiratory failure in the united states, 2008–2012," Annals of the American Thoracic Society, vol. 13, no. 12, pp. 2184–2189, 2016.

[32] G. Shari, M. Kojicic, G. Li, R. Cartin-Ceba, C. T. Alvarez, R. Kashyap, Y. Dong, J. T. Poulose, V. Herasevich, J. A. C. Garza et al., "Timing of the onset of acute respiratory distress syndrome: a population-based study," Respiratory care, vol. 56, no. 5, pp. 576–582, 2011.

[33] E.Apostolova, T.Wang, T.Tschampel,I .Koutroulis, and T.Velez,"Combining structured and free-text electronic medical record data for real-time clinical decision support," in Proceedings of the 18th BioNLP Workshop and Shared Task, 2019, pp. 66–70.

[34] T. Bai, A. K. Chanda, B. L. Egleston, and S. Vucetic, "Ehr phenotyping via jointly embedding medical concepts and words into a unified vector space," BMC medical informatics and decision making, vol. 18, no. 4, p. 123, 2018.

[35] E. Apostolova, A. Uppal, J. Galarraga, I. Koutroulis, T. Tschampel, T. Wang, and T. Velez, "Towards reliable ards clinical decision support: Ards patient analytics with free-text and structured emr data," AMIA Summits on Translational Science Proceedings, vol. 2019, 2019.

[36] Y. Kim, "Convolutional neural networks for sentence classification," arXiv preprint arXiv:1408.5882, 2014.

[37] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," arXiv preprint arXiv:1810.04805, 2018.

AUTHORS

Emilia Apostolova, Ph.D.

Dr. Apostolova has extensive experience in Natural Language Processing and Machine Learning and has developed a number of information extraction and text categorization systems in the biomedical, financial, and legal domains. She has been an NIH Researcher and a Master – Natural Language Processing / Machine Learning at the Deloitte Innovation Lab. Dr. Apostolova received her BA/MA in Linguistics / English Literature from Sofia University St. Kliment Ohridski, MS Computer Science from Loyola University Chicago, and PhD Computer Science / Natural Language Processing from DePaul University.

Joe Morales

Joe Morales began his career as a nuclear submarine office, serving onboard USS Alabama (SSBN-731) where he made six strategic deterrent patrols. After completion of his Navy service, Joe returned home to Colorado Springs and embarked on career in software development. Joe is now a software leader with 25 years experience leading teams in the design, development, and fielding of software systems in a variety of domains for commercial and government customers. He has extensive expertise in data analytics and processing, machine learning, modeling and simulation, and cloud-based architectures. Joe received his BS in Computer Science from the US Naval Academy and Masters of Information and Data Science from the University of California, Berkeley.

Ioannis Koutroulis, MD

Ioannis Koutroulis, M.D., is originally from Greece, where he attended medical school and received a Ph.D. in infectious diseases. He completed residency training in New York and a fellowship in Philadelphia at St. Christopher's Hospital for Children and Drexel University College of Medicine. At that time, he also completed an MBA in healthcare management and finance. While at Drexel, Dr. Koutroulis served as the course director of pathophysiology for second year medical students. He has published work on sepsis, atopic dermatitis, obesity and asthma. Dr. Koutroulis is currently the Co-Director, Urgent Care Fellowship and an Emergency Medicine Physician at Children's National Hospital.

Tom Velez, PhD, JD

Dr. Tom Velez is an accomplished entrepreneur, aerospace engineer, mathematician and computer scientist with over 40 years of experience in the development of complex technology-intensive systems in the aerospace, intelligence, military, and healthcare industries. His current interest and passion is directed at the effective utilization of advanced information systems technologies in support of outcomes-focused, patient-centered chronic disease care quality. Dr. Velez received a B.S. in Mathematics from Iona College, M.S. in Mathematics and Computer Science from Adelphi University; Ph.D. in Applied Mathematics from Georgetown University and J.D., Magna Cum Laude from the University of Baltimore.