Deep multitask ensemble classification of emergency medical call incidents combining multimodal data improves emergency medical dispatch

Pablo Ferri[1], Carlos Sáez[1], Antonio Félix-De Castro[2], Javier Juan-Albarracín[1], Vicent Blanes-Selva[1], Purificación Sánchez-Cuesta[2] and Juan M García-Gómez[1]

[1]Biomedical Data Science Laboratory (BDSLab), Instituto de Aplicaciones de las Tecnologías de la Información y de las Comunicaciones Avanzadas (ITACA), Universitat Politècnica de València (UPV), Valencia, Spain

[2]Conselleria de Sanitat Universal i Salut Pública, Generalitat Valenciana (GVA), Valencia, Spain

Corresponding Author: pabferb2@upv.es

ABSTRACT

Objective: To develop a predictive model to aid non-clinical dispatchers to classify emergency medical call incidents by their life-threatening level (yes/no), admissible response delay (undelayable, minutes, hours, days) and emergency system jurisdiction (emergency system/primary care) in real time.

Materials: A total of 1 244 624 independent retrospective incidents from the Valencian emergency medical dispatch service in Spain from 2009 to 2012, comprising clinical features, demographics, circumstantial factors and free text dispatcher observations.

Methods: A deep multitask ensemble model integrating four subnetworks, composed in turn by multi-layer perceptron modules, bidirectional long short-term memory units and a bidirectional encoding representations from transformers module.

Results: The model showed a F1 score of 0.771 in life-threatening classification, 0.592 in response delay and 0.801 in jurisdiction, obtaining a performance increase of 13.2%, 16.4% and 4.5%, respectively, with regard to the current in-house triage protocol of the Valencian emergency medical dispatch service.

Discussion: The model captures information present in emergency medical calls not considered by the existing in-house triage protocol, but relevant to carry out incident classification. Besides, the results suggest that most of this information is present in the free text dispatcher observations.

Conclusion: To our knowledge, this study presents the development of the first deep learning model undertaking emergency medical call incidents classification. Its adoption in medical dispatch centers would potentially improve emergency dispatch processes, resulting in a positive impact in patient wellbeing and health services sustainability.

Keywords: medical emergencies, emergency medical calls, emergency medical dispatch, deep learning, multitask learning, ensemble learning.

BACKGROUND AND SIGNIFICANCE

Emergency medical dispatch (EMD) involves the reception and management of requests for medical assistance in an emergency medical services system. [1] It comprises two main dimensions: call-taking, where emergency medical calls are received and incidents are classified according to their priority, i.e., triaged; and controlling, where the best available resources are dispatched to handle the event. [2]

The call-taking process is generally managed by emergency medical dispatchers. [3] These mediators are in many cases non-clinical staff, trained with the essential knowledge of medical emergencies for the proper and efficient management of the incident. [1] [4] Dispatchers usually follow a clinical protocol, established in the medical dispatch center, and periodically verified by medical supervisors. [5]

However, despite preparation and the existence of triage protocols, assigning priorities to emergency medical call incidents (EMCI) is a challenging and stressful task for dispatchers, requiring constant concentration. [6] [7] [8] Additionally, there is always an inherent uncertainty on the real patient state, since the information of the event is gathered from telephonic interview processes. Furthermore, there are time constraints due to the incident priority or the need for tackling other incoming calls. [9] A wrong priority assignment derives either in insufficient medical attention or unnecessary resource deployment. [10] [11] [12] In consequence, EMCIs triage protocols are continuously revised and enhanced.

Many triage algorithms, such as the Emergency severity index, [13] the Manchester triage system, [14] the Canadian triage and acuity scale [15] or the Australasian triage scale, [16] have been widely studied and enriched. [17] [18] [19] [20] However, they are difficult to benchmark, deriving in no international agreement about their use for EMD. [21] Likewise, these algorithms depend on structured clinical information which is not always available during the call. [22] As such, improvements in EMD processes by redefining this sort of protocols are extremely costly and limited.

In the Valencian Community (Spain), the triage of EMCI is currently supported by an in-house triage protocol, based on a clinical decision tree, grounded on heavily structured clinical variables, e.g., chest pain (yes or no), collected throughout the interview in a sequential manner. Therefore, free text dispatcher observations, with higher expressiveness than structured data, cannot be automatically processed by the protocol, limiting its generalization to situations beyond the established guidelines.

The potential capability of deep learning to enhance EMCI classification through the provision of decision support to non-clinical dispatchers, was spotted by the Health Services Department of the Valencian region, aware of the potential of these models: deep learning is at the state of the art of machine learning in tasks involving complex types of data, [23] e.g., high dimensional, unstructured, sequential, multimodal, [24] [25] [26] [27] such as those found in EMCI databases. Likewise, this and other machine learning tools have

already been applied to tackle EMD challenges such as ambulance allocation, [28] [29] [30] prediction of emergency calls volume, [31] automatic stress detection of the caller, [32] interpretable knowledge extraction, [33] performance monitoring [34] or cardiac arrest calls assistance. [35] Therefore, we can argue that deep learning models are a feasible and promising technology to improve EMD through EMCI classification.

In this work, we develop and evaluate a deep learning model to provide decision support to non-clinical dispatchers in EMCI triage from the medical dispatch center of the Valencian region. Our model is designed to integrate the EMCI data collected during the call and carry out its classification. To our knowledge, this is the first study undertaking general, i.e., not restricted to specific disorders, EMCI classification trough deep learning.

MATERIALS

Dataset

A total of 1 244 624 independent EMCI of the Health Services Department of the Valencian Community, comprising during-call and after-call data, were compiled in retrospective from 2009 to 2012. The Health Services Department board of the Valencian Community approved the data use for this project, removing before their analysis any information that may disclose the identity of the person.

During-call data (Figure 1 top) consist of demographics, circumstantial factors, clinical features collected throughout the triage tree navigation and free text dispatcher observations. From a data type perspective, we found structured, i.e., fixed fields, and unstructured, i.e., open fields, data, as well as stationary, i.e., their order is not informative, and sequential, i.e., their order is informative, data. Further details about these data are available in Supplementary material Appendix 1.

After-call data involve physician diagnoses standardized by the International classification of diseases codes, [36] maneuvers, procedures, hospitalizations and emergency department stays linked to each one of the incidents (Supplementary material Appendix 1). These data were used to derive EMCI classification labels.

The inclusion criteria in our study consisted in those EMCI which after-call data were fully available, and which during-call data were registered by non-novice dispatchers, i.e., dispatchers with more than 100 calls managed. The final working dataset size comprised 722 270 EMCI.

Labels derivation

Three different but complementary labels were defined to classify EMCI (Figure 1 bottom): life-threatening level (yes/no), admissible response delay (undelayable, minutes, hours, days) and emergency system jurisdiction (emergency system/primary care). These labels were derived from after-call data, by means of a mapping defined by a panel of 17 physicians from the Health Services Department of the Valencian Community, using a Delphi methodology. [37]

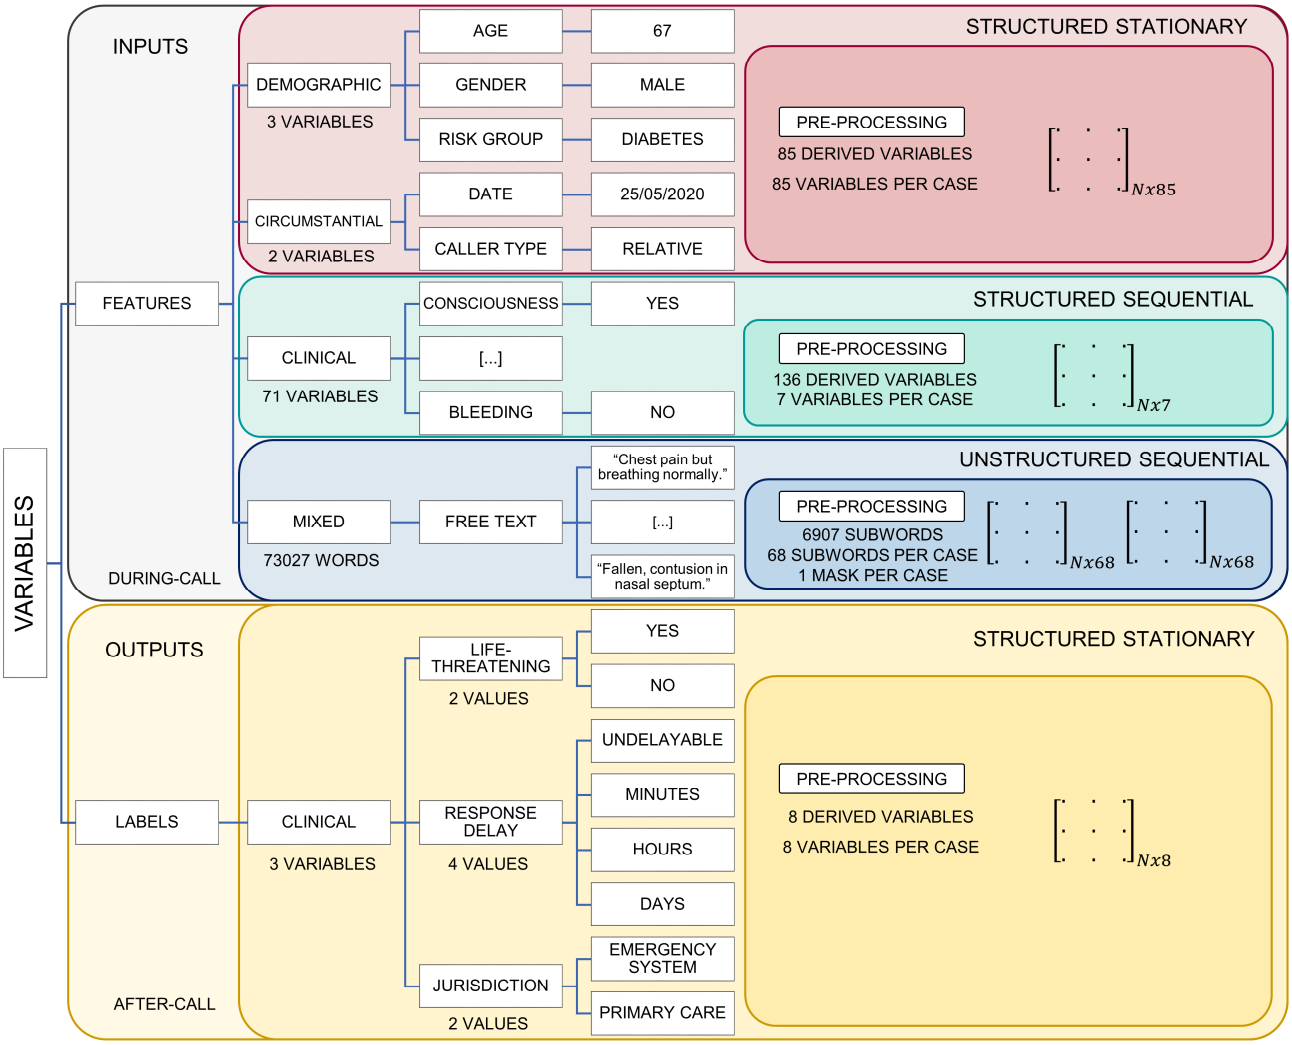

Figure 1. Dataset variables arranged by type. Names and cardinality (before and after pre-processing) are presented, along with examples for their values. N is equal to the 722 270 EMCI used in the study.

Framework

The implementation language was Python 3.7.3, [38] making use of libraries Pandas, [39] NumPy, [40] and Fuzzywuzzy, [41] for data pre-processing and Pytorch (version 1.4.0), [42] Hugginface transformers [43] and Hyperopt [44] for modeling.

METHODS

Data pre-processing

Depending on variable type, different pre-processing techniques were applied, mapping the original data to a matrix representation (Figure 1 right, highlighted pre-processing blocks). Details about this pre-processing step can be found in Supplementary material Appendix 2.

Data splitting and sampling

To evaluate model performance and tune hyperparameters without any bias, data were iteratively and randomly split into six subsets (Figure 2). [45] First, data were randomly split into two disjoint *design* and *test sets*, with 80% and 20% proportions respectively. Next, the *design* set was randomly divided again into a *training* and a *validation set*, with 80% and 20% proportions. Finally, a sampling step was performed taking 100000 elements to define a *training* and a *validation sample*.

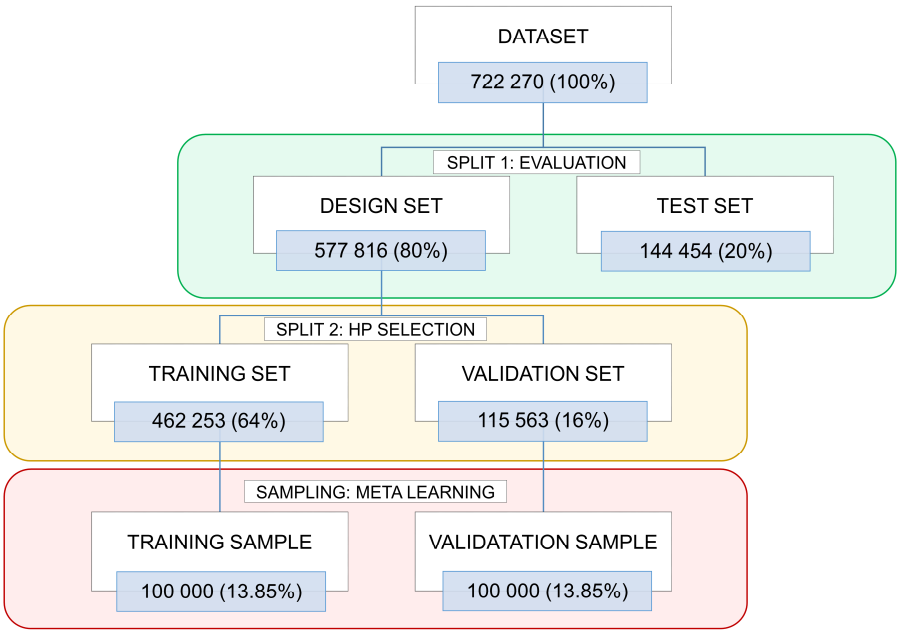

Figure 2. Data splitting and sampling. The number of data of each partition, along with its percentage respect the total number of data, are provided. Abbreviations: HP, hyperparameter.

Deep neural network design

The problem of classifying EMCI combining multimodal data was divided into four subproblems: three EMCI classification problems taking as inputs for each one EMCI data from the same type (structured stationary, structured sequential and unstructured sequential) and a last EMCI classification problem taking as inputs inner outputs obtained from the solution of the prior problems. To solve these four challenges, four deep learning subnetworks were developed: the *Context subnetwork* (ConNet), the *Clinical subnetwork* (CliNet), the *Text subnetwork* (TextNet) and the *Ensemble subnetwork* (EnsNet). Finally, once trained, they were combined in a single global modular neural network model, [46] defining the *Modular network* (ModNet).

Likewise, as the life-threatening, response delay and jurisdiction labels provide different but related information, e.g., a life-threatening situation implies a low admissible response delay, a multitask learning [47] paradigm was followed, to exploit these label dependences. To promote training efficiency and regularization while reducing the number of subnetworks parameters, a hard parameter sharing approach [48] was adopted.

Hence, each of the four developed subnetworks presented a task-shared block (same set of parameters for all label prediction tasks) and a task-specific block (specific set of parameters for each label prediction task).

Every subnetwork is described next, supported by Figure 3, while specific design justifications can be found in Supplementary material Appendix 3:

The *Context subnetwork* (Figure 3 left) deals with the demographics and circumstantial factors (structured and stationary) bound to an EMCI. It consists on a multi-layer perceptron (MLP) [49] composed by dense and output blocks. A dense block integrates a fully connected layer [50] a batch normalization layer [51] a leaky ReLU [52] activation function and a dropout layer, [53] while an output block is composed by a fully connected layer and a softmax activation function.

The *Clinical subnetwork* (Figure 3 center) deals with the clinical features (structured and sequential) collected during the call. It is composed by an embedding layer, [54] a stack of multiple bidirectional long short-term memory (BLSTM) [55] units, multiple skip connections [56] across the BLSTM units, a concatenation block (concatenates the outputs of these skip connections) and a MLP module, integrated by dense and output blocks.

The *Text subnetwork* (Figure 3 right) deals with the free text dispatcher observations (unstructured and sequential) written during an EMCI. It is composed by a bidirectional encoding representations from transformers (BERT) [57] block and a MLP module. The first one is comprised in turn by an embedding block, an encoder block, [58] and a pooler block, while the MLP component is constituted by dense and output blocks.

The *Ensemble subnetwork* (Figure 3 bottom) integrates inner outputs from the ConNet, CliNet and TextNet to generate the final outputs of the *Modular network*. It consists of a concatenation block with a MLP component, composed by dense and output blocks. The inputs of the concatenation block are the outputs of the last layer of the dense block prior to the task-specific block of each one of the former subnetworks.

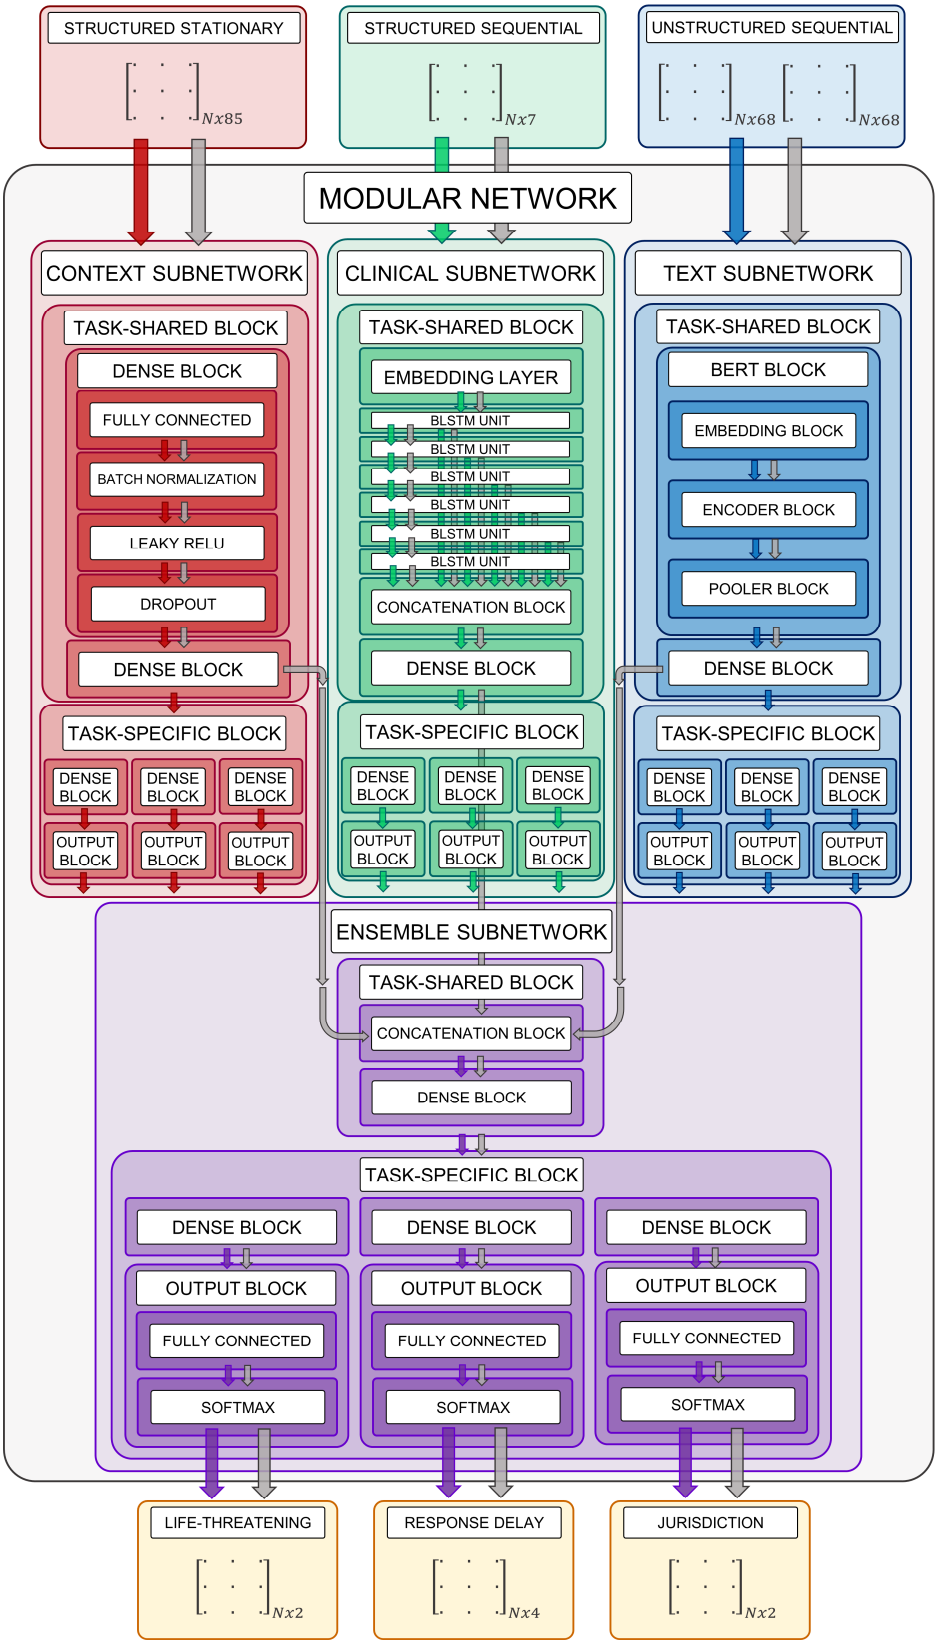

Figure 3. Deep learning model architecture (*Modular network*), including its constituting subnetworks (*Context subnetwork*, *Clinical subnetwork*, *Text subnetwork* and *Ensemble subnetwork*). Arrows indicate the forward propagation direction, for each one of the subnetworks, as well as the global network, colored according to the particular neural network they refer.

Parameter tuning

Subnetworks were trained in a constructive modularized manner, [46] so they were independently trained and assembled later as loosely coupled models. The optimizer selected for that was ADAM, [59] given its learning adaptability, noisy gradients management and learning process stability. [60] [61]. A term of weight decay [62] was included in the parameters upgrading rule expression, to promote regularization. Likewise, it was followed a mini-batch upgrading approach, [63] computing gradients with backpropagation [64] and backpropagation through time. [65] The objective function was a cross-entropy [66] loss (CEL). For each subnetwork, three CEL were calculated (one per label) averaged afterwards and finally backpropagated to carry out the parameter tuning process. Layers with leaky ReLU activation functions were initialized with Kaiming initialization, [67] while softmax activation function layers were initialized with Xavier's initialization. [68]

Hyperparameter tuning

The influence of hyperparameters over subnetworks performance was carefully considered in this work, in order to maximize the attainable outcomes. The hyperparameters studied were related with subnetworks architecture and optimizer settings (check Supplementary material Appendix 4 for details).

Hyperparameters were tuned following a multi-step strategy (Figure 4):

The first step involved an automatic active learning [69] hyperparameter optimization process (Figure 4 top): four surrogate models (one per subnetwork), based on tree-structured parzen estimators, [70] learned the conditional probability distribution of subnetworks hyperparameters given their associated CEL. Aiming to maximize the Expected Improvement [71] of the CEL, new hyperparamter configurations were iteratively sampled from the surrogate models, being upgraded after each training loop. Thereby, 280 different subnetworks (70 hyperparameter configurations times four subnetworks) were trained and evaluated in the *training* and *validation samples*, respectively.

Next, the best hyperparameter configurations proposed by the surrogate models were selected (Figure 4 middle). To prevent overfitting, the best five hyperparameter configurations for each subnetwork were taken to retrain and validate the subnetworks, in the *training* and the *validation set*, respectively, obtaining a total of 20 models trained in this step. Then, the CEL was obtained for each of them and those hyperparameter configurations with the best value, i.e., lowest validation CEL, were considered as the *optimal* hyperparameter configuration.

Finally, the *optimal* hyperparameters were used to retrain the four subnetworks using the whole *design set*, to ensure a proper exploitation of the data (Figure 4 bottom). Once trained, its integration into a single architecture defined the global network (ModNet), evaluated later in the *test set*.

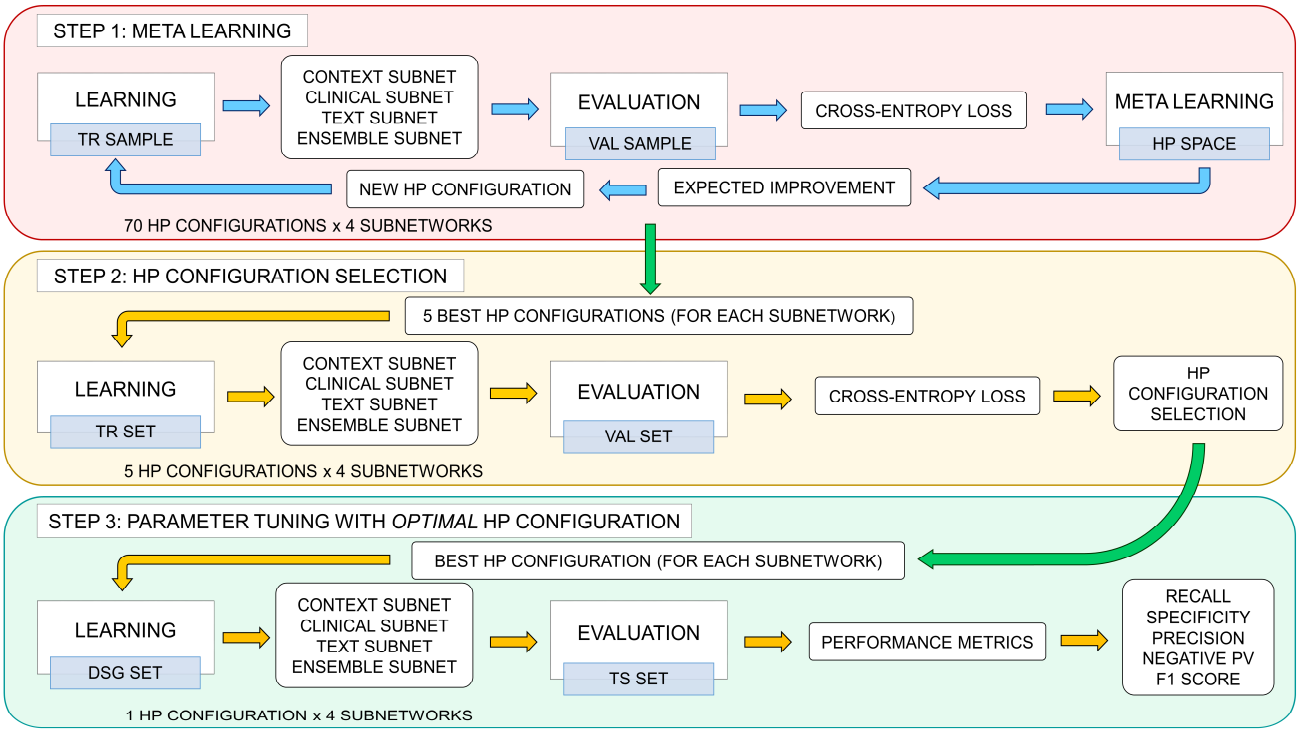

Figure 4. Multi-step hyperparameter tuning strategy. Yellow arrows imply unidirectionality, while blue arrows stand for a feedback loop, both inside a hyperparameter optimization step. Green arrows denote unidirectionality across hyperparameter optimization steps. Abbreviations: HP, hyperparameter; TR, training; VAL, validation; DSG, design TS, test.

Evaluation

The performance of the ModNet, as well as ConNet, CliNet and TextNet subnetworks (EnsNet outputs are the same as the ModNet), were evaluated in the *test set* (144 454 independent EMCI) for each label prediction task. Likewise, performance metrics were also obtained for the current triage protocol of the Valencian emergency medical dispatch service, as a comparative baseline. The evaluation metrics included recall, specificity, precision, negative predictive value (NPV) and the F1 score. For binary labels (life-threatening, jurisdiction), recall, specificity, precision and NPV were referencing the interest class, i.e., life-thread and emergency system jurisdiction. Regarding the multiclass label (response delay), recall, specificity, precision and NPV were calculated for each class and then weighted following a macro approach. Finally, F1 score was computed for all the labels, through a micro weighting schema to dispose of an overall performance descriptor.

RESULTS

Tables 1, 2 and 3 show the classification performance results for the life-threatening, response delay and jurisdiction labels. Metrics are calculated in the *test set*, for the protocol, the ConNet, the CliNet, the TextNet and the ModNet. Percentage differences (Δ) between the ModNet (final deep learning model) and the protocol are also reported.

Table 1. In-house triage protocol and deep learning models performance in life-threatening prediction (test set).

Model	Life-threatening level (yes/no)				
	Recall	Specificity	Precision	Negative PV	F1 score
Protocol	0.644	0.636	0.547	0.723	0.639
ConNet	0.44	0.785	0.583	0.672	0.644
CliNet	**0.79**	0.61	0.581	**0.809**	0.683
TextNet	0.638	**0.844**	0.737	0.773	0.76
ModNet	0.671	0.84	**0.742**	0.789	**0.771**
Δ (%)	2.7	20.4	19.5	6.6	13.2

Abbreviations: PV, predictive value; Δ, ModNet difference respect the protocol.

Table 1 shows that the ModNet outperforms the current protocol in the life-threatening prediction task. It captures more true life-threatening situations (higher recall) being much more precise, i.e., with less false positives (much higher precision). Respect to non-life-threatening incidents, it detects many more true cases of this type (much higher specificity) also with less false negatives (higher NPV). Referring to the overall performance in both classes, the ModNet beats the protocol by far (13.2% of micro F1 score improvement).

Focusing on the subnetworks, the ConNet is the weakest deep learning model, although its F1 is superior to that attained by the protocol. The CliNet offers the better detection rate for true life-threatening situations but at the expense of a significant amount of false positives. Finally, the TextNet exhibits the overall better behavior although its capability to capture true life-threatening events is not the best among the subnetworks.

Table 2 shows that ModNet outcomes are substantially above those achieved by the protocol in the response delay prediction task. Overall detection of situations with a specific true admissible response delay (undelayable, minutes, hours, days) is amply improved by the ModNet (15.8% increment in macro recall) while remarkably enhancing overall precision (17.3% raise). Regarding to the overall capturing of events which do not exhibit certain true admissible response delay, the ModNet is also superior (5.8% increase in macro

specificity) showing less false negatives in this task (5.5% increment in macro NPV). Concerning to the general performance in all classes, the ModNet significantly exceeds the protocol (16.4% of micro F1 score improvement).

Table 2. In-house triage protocol and deep learning models performance in response delay prediction (test set).

Model	Admissible response delay (undelayable, minutes, hours, days)				
	Recall	Specificity	Precision	Negative PV	F1 score
Protocol	0.411	0.8	0.416	0.805	0.428
ConNet	0.376	0.791	0.415	0.793	0.413
CliNet	0.477	0.824	0.53	0.829	0.506
TextNet	0.544	0.851	0.583	0.854	0.576
ModNet	**0.569**	**0.858**	**0.589**	**0.86**	**0.592**
Δ (%)	15.8	5.8	17.3	5.5	16.4

Abbreviations: PV, predictive value; Δ, ModNet difference respect the protocol.

Focusing on ModNet subnetworks for response delay prediction, the ConNet is at the bottom in performance terms, not being capable of outperforming the protocol. The CliNet is clearly over the ConNet and already beats the protocol, while the TextNet is the best ModNet subnetwork in all metrics, with a substantial increase respect to the CliNet.

Table 3. In-house triage protocol and deep learning models performance in jurisdiction prediction (test set).

Model	Emergency system jurisdiction (emergency system/primary care)				
	Rec	Spec	Prec	Negative PV	F1
Protocol	0.855	0.541	0.8	0.635	0.756
ConNet	**0.945**	0.288	0.741	0.708	0.736
CliNet	0.9	0.521	0.802	0.708	0.78
TextNet	0.917	0.519	0.804	**0.745**	0.791
ModNet	0.895	**0.597**	**0.827**	0.726	**0.801**
Δ (%)	4	5.6	2.7	9.1	4.5

Abbreviations: PV, predictive value; Δ, ModNet difference respect the protocol.

Table 3 shows that the ModNet outperforms the protocol in the jurisdiction prediction task. It captures more situations which are truly jurisdiction of the emergency system (better recall) being more precise, i.e., with less false positives (better precision). In relation with incidents which should be derived to primary care, i.e., non-emergencies, the ModNet detects more true cases of this type (higher specificity) also with less false negatives (better NPV). Respect to the overall performance in both classes, the ModNet surpasses the protocol (4.5% of micro F1 score improvement).

Regarding to ModNet subnetworks, although the ConNet presents the highest recall values, its specificity is fairly poor, with worse general results than the protocol in the jurisdiction prediction task. The CliNet provides a substantial improvement over the later subnetwork, with an overall performance above the protocol. As in life-threatening and response delay, the TextNet is the subnetwork attaining the best outcomes.

DISCUSSION

Relevance

The superior performance of the ModNet against the triage protocol suggests the existence of information provided during the emergency medical call not considered by the current protocol, but captured by the deep learning model. According to TextNet outcomes, far better than those attained by the ConNet and CliNet, most of this information would be present in the free text dispatcher observations. Since text fields are unbounded, they would embrace wider casuistry, allowing more precision in the EMCI description, lowering, consequently, its uncertainty.

Clinical variables stand as excellent life-threatening detection (about 80% of total cases) features, since dispatchers ask them to reduce chances of missing situations where patient's life is at risk. Likewise, the outstanding emergency system jurisdiction recall of demographics and circumstantial factors (capturing about 95% of total cases) may be related with patient profiles highly susceptible from requiring emergency aid, e.g., elderly cardiac patient males.

The hardest classification problem is to predict the admissible response delay, probably derived from the fact that it is a multiclass label, presenting twice possible outputs (undelayable, minutes, hours, days) than the other labels (life-threatening, jurisdiction), which are binary.

The modular approach followed in this work, assembling four specialized subnetworks into a single global network, has shown that the potential of the aggregated network is superior to any of its individual components, balancing their respective weaknesses and strengths while properly integrating processed information within each one.

Finally, the results of this work imply that current emergency dispatch processes could be improved by means of deep learning, eventually deriving in a positive impact over patient wellbeing and health services sustainability.

Limitations

The main limitation of this work is the inherent uncertainty of the problem: in the studied dataset it was likely to find rather similar input combinations presenting completely different label values. In other words, different disorders presented the same clinical picture. For example, chest pain may imply a life-threatening situation, if the underlying unknown cause is a heart attack, or not, since it could be derived from a prior anxiety crisis. This non-discriminative variability sets bounds in terms of maximum performance attainable by any model, i.e., Bayes error. [72]

Future work

Next steps include the evaluation of the deep learning model with prospective cases from the Valencia region and its deployment and integration in the emergency medical dispatch center. For that, we will propose a graphical user interface to allow the interaction between the dispatcher and the model during the call. Finally, the resulting tool will be implemented in the emergency medical dispatch center of the Valencian Community.

CONCLUSIONS

A novel deep multitask ensemble model, designed to aid non-clinical dispatchers during emergency medical calls to classify incidents by their life-threatening level, admissible response delay and emergency system jurisdiction, has been developed and successfully evaluated. To our knowledge, this is the first deep learning model implemented to face this challenge.

The performance achieved by the model is notably superior to that attained by the current in-house triage protocol of the emergency medical dispatch service of the Valencian Community, achieving an improvement of 13.2%, 16.4%, 4.5% in life-threatening, response delay and jurisdiction classification, respectively, with regard to the F1 score metric.

The network modular design with specialized subnetworks for the different data modalities has allowed discovering the potential benefit of the information contained in free text fields for the automatic classification of emergency medical call incidents. This information can be used to optimize current guidelines.

The implantation of this model in medical dispatch centers would have a remarkable impact in patient wellbeing and health services sustainability.

FUNDING

This work has been supported by the *Valencian agency for security and emergency response* project A180017304-1, the *Ministry of Science, Innovation and Universities* of Spain program FPU18/06441 and the *EU Horizon 2020* project InAdvance 825750.

AUTHOR CONTRIBUTIONS

AFC and PSC led the definition of medical requirements of the study, exported data from database, conducted data labeling and review. PF and CS pre-processed data. PF, JJA, VBS and CS developed the deep learning model. PF, CS, AFC and JMG conceived the design of the solution and written the main manuscript. All the authors contributed to the review and revisions of manuscript, and have seen and approved the final version of the manuscript.

SUPPLEMENTARY MATERIAL

Supplementary material is available at *medRxiv* online.

ACKNOWLEDGEMENTS

We thank the support of physicians from the Health Services Department of Valencian Community. CS acknowledges the support of NVIDIA GPU Grant Program.

REFERENCES

[1] J. J. Clawson and K. B. Dernocoeur, "Principles of emergency medical dispatch," *Priority Press,* 2003.

[2] A. Blandford and B. W. Wong, "Situation awareness in emergency medical dispatch," *International journal of human-computer studies 61.4,* pp. 421-452, 2004.

[3] S. J. Stratton, "Triage by emergency medical dispatchers," *Prehospital and disaster medicine,* pp. 263-269, 1992.

[4] J. J. Clawson, "Dispatch priority training: strengthening the weak link," *JEMS,* pp. 32-35 (6), 1981.

[5] L. Palumbo, J. Kubincanek, C. Emerman, N. Jouriles, R. Cydulka and B. Shade, "Performance of a system to determine EMS dispatch priorities," *The American journal of emergency medicine,* pp. 388-390 (14), 1996.

[6] L. Weibel, I. Gabrion, M. Aussedat and G. Kreutz, "Work-related stress in an emergency medical dispatch center," *Annals of emergency medicine,* pp. 500-506 (41), 2003.

[7] K. Forslund, A. Kihlgren and M. Kihlgren, "Operators' experiences of emergency calls.," *Journal of telemedicine and telecare,* pp. 290-297 (10), 2004.

[8] B. Ek, P. Edström, A. Toutin and M. Svedlund, "Reliability of a Swedish pre-hospital dispatch," *International emergency nursing,* pp. 143-149, 2013.

[9] J. Leprohon and V. L. Patel, "Decision-making strategies for telephone triage in emergency medical services," *Medical Decision Making,* pp. 240-253 (15), 1995.

[10] M. Srámek, W. Post and R. W. Koster, "Telephone triage of cardiac emergency calls by dispatchers: a prospective study of 1386 emergency calls," *Heart,* pp. 440-445 (71), 1994.

[11] Institute of Medicine Committee on the Future of Emergency Care in the US Health System, "Emergency medical services: at the crossroads," *Washington: DC,* 2006.

[12] L. Hjälte, B.-O. Suserud, J. Herlitz and I. Karlberg, "Why are people without medical needs transported by ambulance? A study of indications for pre-hospital care," *European Journal of Emergency Medicine,* pp. 151-156 (14), 2007.

[13] Agency for healthcare research and quality, "Emergency severity index implementation handbook," 2012.

[14] Manchester triage group, "Emergency triage," 2014.

[15] Échelle de triage et de gravité, "Revisions to the Canadian emergency department triage and acuity scale (CTAS). Guidelines 2016," 2016.

[16] Australasian College for Emergency Medicine, "A National Triage Scale for Australian Emergency Departments [position paper]," *Melbourne, Victoria, Australia: Australasian College for Emergency Medicine,* 1993.

[17] M. Christ, F. Grossmann, D. Winter, R. Bingisser and E. Platz, "Modern triage in the emergency department," *Deutsches Ärzteblatt International,* p. 892 (107), 2010.

[18] M. N. Storm-Versloot, D. T. Ubbink, J. Kappelhof and J. S. Luitse, "Comparison of an informally structured triage system, the emergency severity index, and the manchester triage system to distinguish patient priority in the emergency department," *Academic Emergency Medicine,* pp. 822-829 (18), 2011.

[19] N. Seiger, M. van Veen, E. W. Steyerberg, M. Ruige, A. H. Van Meurs and H. A. Moll, "Undertriage in the Manchester triage system: an assessment of severity and options for improvement," *Archives of disease in childhood,* pp. 653-657 (96), 2011.

[20] J. M. Zachariasse, N. Seiger, P. P. Rood, C. F. Alves, P. Freitas, F. J. Smit, G. R. Roukema and H. A. Moll, "Validity of the Manchester Triage System in emergency care: A prospective observational study," *PloS one,* 2017.

[21] G. FitzGerald, G. A. Jelinek, D. Scott and M. F. Gerdtz, "Emergency department triage revisited," *Emergency Medicine Journal,* pp. 86-92 (27), 2010.

[22] L. Farand, J. Leprohon, M. Kalina, F. Champagne, A. P. Contandriopoulos and A. Preker, "The role of protocols and professional judgement in emergency medical dispatching," *European journal of emergency medicine: official journal of the European Society for Emergency Medicine,* pp. 136-148 (2), 1995.

[23] Y. LeCun, Y. Bengio and G. Hinton, "Deep learning," *Nature,* pp. 436-444 (521) , 2015.

[24] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-R. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath and B. Kingsbury, "Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups," *IEEE Signal processing magazine,* pp. 82-97 (29), 2012.

[25] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg and L. Fei-Fei, "Imagenet large scale visual recognition challenge," *International journal of computer vision,* pp. 211-252, 2015.

[26] J. Hirschberg and C. D. Manning, "Advances in natural language processing," *Science,* pp. 261-266 (349), 2015.

[27] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel and D. Hassabis, "Mastering the game of Go with deep neural networks and tree search," *nature,* p. 484, 2016.

[28] M. S. Maxwell, S. G. Henderson and H. Topaloglu, "Ambulance redeployment: An approximate dynamic programming approach," *Winter Simulation Conference,* pp. 1850-1860, 2009.

[29] L. A. McLay and M. E. Mayorga, "A model for optimally dispatching ambulances to emergency calls with classification errors in patient priorities," *IIE Transactions,* pp. 1-24 (45), 2013.

[30] A. Y. Chen and T.-Y. Lu, "A GIS-based demand forecast using machine learning for emergency medical services," *Computing in Civil and Building Engineering,* pp. 1634-1641, 2014.

[31] N. Channouf, P. L'Ecuyer, A. Ingolfsson and A. N. Avramidis, "The application of forecasting techniques to modeling emergency medical system calls in Calgary, Alberta," *Health care management science,* pp. 25-45 (10), 2007.

[32] I. Lefter, L. J. Rothkrantz, D. A. van Leeuwen and P. Wiggers, "Automatic stress detection in emergency(telephone) calls," *International Journal of Intelligent Defence Support Systems,* pp. 148-168 (4), 2011.

[33] F. Barrientos and G. Sainz, "Interpretable knowledge extraction from emergency call data based on fuzzy unsupervised decision tree," *Knowledge-based systems,* pp. 77-87 (25), 2012.

[34] P. Klement and V. Snásel, "Using SOM in the performance monitoring of the emergency," *Simulation Modelling Practice and Theory,* pp. 98-109 (19), 2011.

[35] S. N. Blomberg, F. Folke, A. K. Ersboll, H. C. Christensen, C. Torp-Pedersen, M. R. Sayre, C. R. Counts and K. L. Freddy, "Machine learning as a supportive tool to recognize cardiac arrest in emergency calls," *Resuscitation,* pp. 322-329, 2019.

[36] World health organization, "International classification of diseases, ICD-9," 2015.

[37] N. C. Dalkey, "The Delphi method: an experimental study of group opinion," *RAND CORP SANTA MONICA CALIF,* 1969.

[38] Python Software Foundation, "Python Language Reference, version 3.6.5," *python.org,* 2018.

[39] W. McKinney, "Data structures for statistical computing in python," *Proceedings of the 9th Python in Science Conference,* pp. 51-56 (445), 2010.

[40] S. v. d. Walt, S. C. Colbert and G. Varoquaux, "The NumPy Array: A Structure for Efficient Numerical Computation," *Computing in Science & Engineering,* pp. 22-30 (13), 2011.

[41] J. Gonzalez, P. Rodrigues and A. Cohen, "Fuzzywuzzy: Fuzzy string matching in python," 2017.

[42] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. De Vito, Z. Lin, A. Desmaison, L. Antiga and A. Lerer, "Automatic differentiation in pytorch," *31st Conference on Neural Information Processing Systems (NIPS),* 2017.

[43] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz and J. Brew, "HuggingFace's Transformers: State-of-the-art Natural Language Processing,"

Huggingface's transformers: State-of-the-art natural language processing, pp. ArXiv, abs/1910.03771, 2019.

[44] J. Bergstra, B. Komer, C. Eliasmith, D. Yamins and D. D. Cox, "Hyperopt: a python library for model selection and hyperparameter optimization," *ICML Workshop on Automated Machine Learning,* p. (9), 2014.

[45] R. Kohavi, "A study of cross-validation and bootstrap for accuracy estimation and model selection," *Ijcai,* pp. 1137-1145 (14), 1995.

[46] K. Chen, "Deep and modular neural networks," *Handbook of Computational Intelligence,* pp. 473-494, 2015.

[47] R. Caruana, "Multitask learning," *Machine learning,* pp. 41-75 (28), 1997.

[48] S. Ruder, "An Overview of Multi-Task Learning in Deep Neural Networks," *arXiv preprint arXiv:1706.05098,* 2017.

[49] F. Rosenblatt, "Principles of neurodynamics. perceptrons and the theory of brain mechanisms," *Cornell Aeronautical Lab Inc Buffalo NY,* pp. No. VG-1196-G-8, 1961.

[50] I. Goodfellow, Y. Bengio and A. Courville, "Deep learning," *MIT press,* 2016.

[51] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," *arXiv preprint arXiv:1502.03167 ,* 2015.

[52] A. L. Maas, A. Y. Hannun and A. Y. Ng, "Rectifier nonlinearities improve neural network acoustic models," *Proc. icml,* 2013.

[53] G. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever and R. Salakhutdinov, "Improving neural networks by preventing co-adaptation of feature detectors," *arXiv preprint arXiv:1207.0580,* 2012.

[54] Y. Bengio, R. Ducharme, V. Pascal and C. Jauvin, "A neural probabilistic language model," *Journal of machine learning research,* pp. 1137-1155 (3), 2003.

[55] M. Schuster and P. K. Kuldip, "Bidirectional recurrent neural networks," *IEEE Transactions on Signal Processing,* pp. 2673-2681 (45), 1997.

[56] K. He, X. Zhang, S. Ren and J. Sun, "Deep residual learning for image recognition," *Proceedings of the IEEE conference on computer vision and pattern recognition,* pp. 770-778, 2016.

[57] 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.

[58] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser and I. Polosukhin, "Attention is all your need," *Advances in neural information processing systems,* pp. 5998-6008, 2017.

[59] D. P. Kingma and J. L. Ba, "Adam: A method for stochastic optimization," *International Conference on Learning Representations (ICLR),* 2015.

[60] S. Ruder, "An overview of gradient descent optimization algorithms," *arXiv preprint arXiv:1609.04747,* 2016.

[61] S. Sun, Z. Cao, H. Zhu and J. Zhao, "A Survey of Optimization Methods from a Machine Learning Perspective," *IEEE transactions on cybernetics,* 2019.

[62] A. Krogh and J. A. Hertz, "A simple weigth decay can improve generalization," *Advances in neural information processing systems,* pp. 950-957, 1992.

[63] D. P. Bertsekas, "Incremental least squares methods and the extended Kalman filter," *SIAM Journal on Optimization,* pp. 807-822 (6), 1996.

[64] R. Hecht-Nielsen, «Theory of the backpropagation neural network,» *Neural networks for perception. Academic Press,* pp. 65-93, 1992.

[65] P. J. Werbos, "Backpropagation through time: what it does and how to do it," *Proceedings of the IEEE ,* pp. 1550-1560, 1990.

[66] K. Janocha and W. M. Czarnecki, "On Loss Functions for Deep Neural Networks in Classification," *arXiv preprint arXiv:1702.05659,* 2017.

[67] K. He, X. Zhang, S. Ren and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," *Proceedings of the IEEE international conference on computer vision,* pp. 1026-1034, 2015.

[68] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," *Proceedings of the thirteenth international conference on artificial intelligence and statistics,* pp. 249-256, 2010.

[69] B. Settles, "Active learning literature survey," *University of Wisconsin-Madison Department of Computer Sciences,* 2009.

[70] J. Bergstra, R. Bardenet, Y. Bengio and B. Kégl, "Algorithms for hyper-parameter optimization," *Advances in neural information processing systems,* pp. 2546-2554, 2011.

[71] D. R. Jones, "A taxonomy of global optimization methods based on response surfaces," *Journal of global optimization,* pp. 345-383 (21), 2001.

[72] K. Fukunaga, "Introduction to statistical pattern recognition," *Elsevier,* 2013.