Joint Prediction of Cocaine Craving and Euphoria using Structured Prediction Energy Networks

Bhanu Teja Gullapalli
University of Massachusetts Amherst
Amherst, MA, USA
bgullapalli@umass.edu

Gustavo A Angarita
Yale University School of Medicine
New Haven, CT, USA
gustavo.angarita@yale.edu

Deepak Ganesan
University of Massachusetts Amherst
Amherst, MA, USA
dganesan@cs.umass.edu

Tauhidur Rahman
University of Massachusetts Amherst
Amherst, MA, USA
trahman@cs.umass.edu

ABSTRACT

In recent years, wearable and mobile health sensing technologies have been developed to track drug usage and monitor different addiction-related states, including craving and euphoria. These states are interdependent and correlated, which is well documented in the literature. However, the state of the art digital biomarker technologies model these states independent of each other and thus fail to use the inherent relationship while making predictions. In our current work, we demonstrate how structured prediction energy networks (SPENs) can be used to capture the correlation and dependencies between self-reported craving, euphoria, and the underlying physiological biomarkers. More specifically, we use SPENs to jointly predict self-reported visual analog scale (VAS) ratings of cocaine craving and euphoria from cardiac signals captured from a wearable chest band. The proposed SPEN-based model can improve the performance of both VAS craving and VAS euphoria prediction by a Normalized Root Mean Square Error of respectively 4.6% and 5.4%.

CCS CONCEPTS

• **Human-centered computing** → **Ubiquitous and mobile devices**.

DigiBiom'21, June 25, 2021, Virtual, WI, USA

ACM ISBN 978-1-4503-8606-7/21/06…$15.00
https://doi.org/10.1145/3469266.3469881

KEYWORDS

Craving, Euphoria, Drug-seeking behavior, On-body sensing, Cardiac signal, Energy networks

ACM Reference Format:
Bhanu Teja Gullapalli, Gustavo A Angarita, Deepak Ganesan, and Tauhidur Rahman. 2021. Joint Prediction of Cocaine Craving and Euphoria using Structured Prediction Energy Networks. In *Future of Digital Biomarkers (DigiBiom'21), June 25, 2021, Virtual, WI, USA.* ACM, New York, NY, USA, Article 111, 7 pages. https://doi.org/10.1145/3469266.3469881

1 INTRODUCTION

Drug addiction is a brain-based chronic disorder that affects a person's behavior and leads to an inability to control drug usage. To treat drug addiction, it is essential to understand the real-world antecedents of recurrent drug usage. For a person addicted to the drug, the cycle of addiction happens as shown in figure 1. Addiction starts due to various reasons, including exposure to any drug-related cues or stressors or a priming dose of the drug, which then transforms into a craving. Eventually, after a loss of control, the individual seeks the drug and finally administers it. After administrating the drug, the individual then enters a state of euphoria, which we commonly call "HIGH". With each euphoric event, the individual develops a new tolerance to the drug, referred to as neuroadaptation. Sometimes the person may administer the drug without any drug-related cues. Recent advancements in ubiquitous and mobile health technologies were able to use various on-body or contactless sensors to predict different variables of the addiction loop for different drugs. However, currently, no system explicitly considers these interactions between the predicted classes while making predictions. In our current work, we present a Structured energy prediction network (SPEN) [3] which jointly predicts craving and euphoria scores of cocaine using cardiac signals obtained from a wearable chest band along with capturing the dependencies

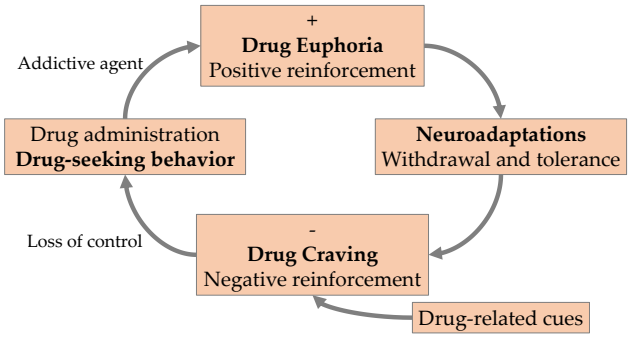

Figure 1: The addiction loop [10].

between the labels.

Previous work by [11, 20] built individual systems for detecting drug administration, predicting craving and euphoria levels of cocaine. These systems do not leverage the information that drug craving, drug administration, and drug euphoria are inter-related to each other (figure 1), but rather make their predictions independent of each other. Essentially, in the cocaine dataset collected by [11], we observed that the ground truth craving and euphoria scores are positively correlated ((rs(121) = 0.64, $p < .001$)). This means that if the person is not feeling any craving for the drug, then the person is not experiencing any euphoric feeling, and if the person is feeling euphoric due to the drug, then the person would like to continue taking the drug by showing a craving to keep experiencing the ongoing effect of the drug. We leverage this information in our SPEN architecture to jointly predict craving and euphoria scores of cocaine. In summary, our contributions are-

- Using SPEN architecture on the previously collected dataset by [11], we show that we can improve the performance of VAS craving prediction from 18.6% Normalized Root Mean Square Error (NRMSE)to 14.2% NRMSE and VAS euphoria prediction from 19.9% NRMSE to 14.5% NRMSE.
- We show that for the participants who do not show any correlation between ground truth VAS craving and euphoria scores, the SPEN architecture performed the same or worse compared to the [11] architecture. In other words, all the participants who showed a correlation between the labels had an improvement in the performance with our architecture.

2 RELATED WORK

Various works had tried to understand and predict different states of the addiction loop using mobile health and ubiquitous technologies for a different types of drugs. Most of these works focussed on detection of drug administration: Cocaine [12, 19, 20], Opioid [5, 17], Alcohol [4, 22], Tobacco [14, 23] as it is one of the initial and central question

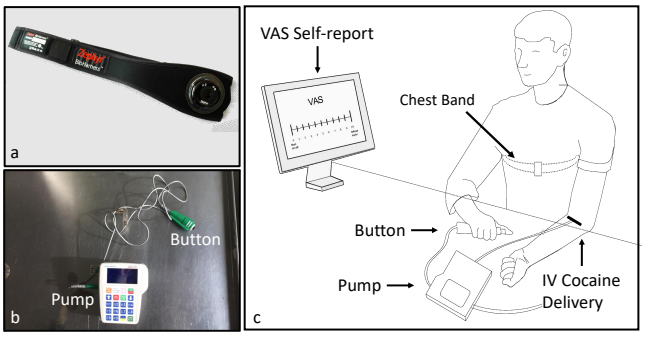

Figure 2: a) Zephyr chest band. b) Infusion Pump button, used to request cocaine in Self-administration period. c) Brief sketch of participant during the Self-administration period of the study.

to ask if one wants to treat drug addiction. With regards to other drug-induced states, Gullapalli et al. [11] used cardiac and respiratory signals from a wearable chest band to detect craving and euphoria. Similarly, detecting cravings of alcohol [2, 24] and smoking Chatterjee et al. [7] was done by using mobile and physiological data collected from wearable wrist band. All these works built systems that address a single part of the addiction loop. On the contrary, in our current work, we show that our system can jointly predict different states of the addiction loop (craving and euphoria) while using the correlation existing between them.

Multilabel predictions, which exploit the correlation between labels, have been studied for a long time, especially in the fields of computer vision [9, 15], natural language processing [8, 26], and bioinformatics [13, 16], etc. These works can broadly be classified into two types- 1) Assume a prior distribution on the labels and use this information while making predictions, 2) Use a machine learning-based architecture on top of feature representations to capture the label dependencies. The latter of these two types has been widely used in recent works as large labeled datasets can be collected relatively easily in the fields mentioned above. In our current work, we use a Structured energy prediction network (SPEN) Belanger and McCallum [3], a machine learning-based architecture to jointly predict cocaine craving and euphoria scores using cardiac signal while exploiting the dependencies between these scores. The primary reason for choosing SPENs over other existing architecture is that it makes no assumptions on the dataset and is proven to work well on small datasets like as in our case.

3 USER STUDY

The following study was conducted at Yale University School of Medicine under the National Institute on Drug Abuse (NIDA) funded research. This study was IRB approved and

shown previously to be safe, well-tolerated, valid, behaviorally relevant, and test-retest reliable [25], and was conducted in the presence of a study physician, advanced cardiac life support certified research nurse, and a basic life support research assistant.

A total of 10 participants were recruited for this study, and each participant was wearing a *Zephyr Bioharness 3* chest band [1] for the entire duration of the study. The chest band captures cardiac and respiratory signals in a passive, continuous, and relatively unobtrusive manner. We only focus on the cardiac signal for our current work due to the poor quality of the respiratory signal captured by the chest band. The entire study for each participant lasted approximately 6 hours, and throughout this study, for every five minutes, the participants were asked to self-report cocaine craving and euphoria according to a visual analog scale (VAS) between 0 ("not at all") to 10 ("most ever"). Below please find the list of questions, their scale, and their acronym. In order to refer to the craving or euphoria/high self-reports, we will use the terms respectively **VAS Craving** and **VAS Euphoria** throughout the paper.

- **VAS Craving** (scale 0-10): how much are you *craving for cocaine* now?
- **VAS Euphoria** (scale 0-10): how *high or euphoric* are you feeling now?

The study comprised of three distinct periods, in the following order: a) an initial drug-free baseline period, b) a subsequent fixed-order, escalating dose, bolus cocaine administration period, and c) a final self-regulated/administrated, ad-libitum (i.e., "binge") cocaine administration period.

Baseline and Fixed-order periods: The study starts with a 30-minute baseline period during which the participant did not receive any cocaine. The baseline period is followed by the fixed dosage period, where three separate bolus intravenous (IV)- 8, 16, and 32 mg IV per 70kg body weight (with a 100 kg maximum cap) are administered at a 20-minute interval. The primary purpose of these periods is to ensure that the participant does not show any unusual behavioral or cardiovascular effects due to the cocaine administration.

Self-administration period: This period was designed to simulate a period of self-regulated, "binge" cocaine consumption. During this period, a participant obtained cocaine via self-initiated presses of a corded infusion pump button (shown in figure 2). The minimum interval between successive cocaine administrations is kept five minutes to ensure the safety of the participants. The self-administration period consists of three one-hour sessions, during which subjects received each of the three cocaine dosage types (i.e., 8mg, 16mg, or 32mg/70kg IV) under a fully randomized, double-blind schedule. During a given 1 hr self-administration session, only a single dosage type is available to the subject and would receive this amount for each bolus in that session. Figure 2 shows the *Zephyr* chest band sensor used, infusion pump button, and finally, a brief sketch of the participant during the study and figure 3 shows an illustration of different periods of the study and self-reports for craving and euphoria for one of the participants. For detailed user study design, please refer [11].

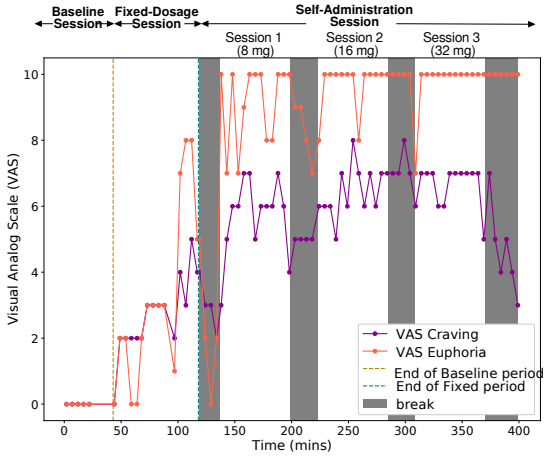

Figure 3: Illustration of VAS craving and euphoria scores for one of the participants. In this case, as the participant starts feeling euphoric, the craving intensity increases.

4 MODEL ARCHIECTURE

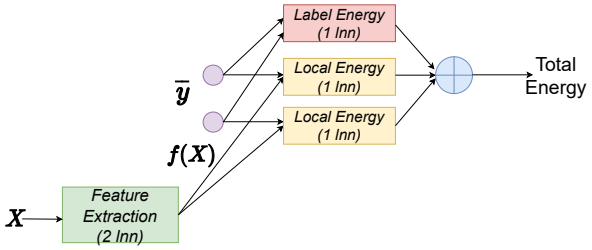

Figure 4: SPEN architecture for multilabel (craving and euphoria VAS) prediction

Throughout the study, for every 5 minutes, the participants had to self-report their cocaine craving and euphoria scores. Results of the Spearman correlation between these scores collected from all participants indicated that there was a significant positive association ($rs(121) = 0.64$, $p < .001$). Therefore, a system designed to predict both craving and euphoria scores jointly should consider this correlation along with the input from the cardiac signal. To achieve this, we decided to use SPEN [3] architecture. While many other similar architectures consider dependencies among labels to make predictions, SPEN requires no assumption of the labels' distribution and works very well with small datasets. Our model

takes features extracted from every 5-minute window as an input to predict VAS craving and euphoria scores jointly. The input features we considered include statistical features of the morphological descriptors of the ECG signal. The statistical features we used are Minimum, Maximum, Median, Mean, Standard deviation, 33 percentile, 67 percentile, Skewness, and Kurtosis. The set of features we considered is the same as the one used in the previous work [11].

SPEN: Let the space of input be $\mathcal{X}$ and output space $\mathcal{Y}$=$\{0, 1, ..., 10\}^2$ as here we have two labels and VAS reported is an integer value in 0-10 range. The core idea behind SPEN is an energy function E_θ:$\mathcal{X} \times \mathcal{Y} \mapsto \mathbb{R}$, that uses a machine learning-based architecture parameterized by θ to produce a continuous energy value for every input/output pair. The energy function E is defined by the sum of two terms local energy- E_{local} and label energy- E_{label}. Given an input x and an output $\bar{y}$, the local energy is expressed as the sum of linear models over the label space given by the equation:

$$E_{local}(x, \bar{y}) = \sum_{l=1}^{2} \bar{y}_l b_l^T f(x)$$

where $\bar{y}_l$ is output for label l (craving or euphoria), b_l is the parameter vector corresponding to label l and $f(x)$ is the representation generated for input x by our feature extraction module. We use a simple 2-layer neural network for feature extraction, similar to the original SPEN architecture, and pre-train it beforehand for convenience. Next, for the label energy, its primary purpose is to capture all label interactions independent of x. The equation is given by:

$$E_{label}(\bar{y}) = c_2^T g(C_1 \bar{y})$$

The product $C_1\bar{y}$ captures the features from the labels, which in turn are used to model their dependencies, g is a non-linear operation which in our case is ReLU[18]. Total energy E_θ therefore is given by :

$$E_\theta(x, \bar{y}) = E_{local}(x, \bar{y}) + E_{label}(\bar{y})$$

The SPEN model architecture for jointly predicting VAS Craving and VAS Euphoria scores is shown in Figure 4.

Learning SPEN: Learning the parameters of SPEN is a two-step process. First, we want to learn the energy function E_θ parameterized by θ. Second, once the energy function is learned, we want to find the $\hat{y}$, which results in minimum energy.

$$\hat{y} = \min_{\bar{y}} E_\theta(x, \bar{y}) \quad (1)$$

To train the energy network we use a structured loss similar to [3]:

$$\min_\theta \sum_{(x_i, y_i)} \max_{\bar{y}_i} [\Delta(y_i, \bar{y}_i) - E_\theta(x_i, \bar{y}_i) + E_\theta(x_i, y_i)]_+ \quad (2)$$

(x_i, y_i) are all the ground truth input/output pairs from the training dataset, $[f]_+$=max(0,f). Traditionally, to minimize this loss, we perform stochastic gradient descent with respect to the energy parameters θ. However in Equation 2, for every (x_i, y_i) pair we would have to compute $\text{argmax}_{\bar{y}_i}(\Delta(y_i, \bar{y}_i) - E_\theta(x_i, \bar{y}_i))$ to find the $\bar{y}_i$ which results in the minimum loss . This means every time we update the energy parameters θ during the gradient descent step, for every data point, we would have to consider all combinatorial possibilities for $\bar{y}_i$ to calculate the loss, making this process very expensive. Therefore, [3] perform gradient descent on $\bar{y}_i$. Initially all the energy parameters θ and $\bar{y}_i$ are randomly initialized, and every time we perform gradient descent to minimize the loss we update both these parameters. During the inference time, the parameters of the energy network are fixed while we perform gradient descent on the outputs to iteratively optimize the energy function and find the optimal $\bar{y}_i$ for every data point. $\Delta(y_i, \bar{y}_i)$ is the error function between ground truth and predicted labels, while [3] used structured hinge loss as they were doing multilabel classification we decided to use a root-mean-squared-error as our output space is $\mathcal{Y}$=$\{0, 1, ..., 10\}^2$ and predictions are continuous values. $E_\theta(x_i, y_i) - E_\theta(x_i, \bar{y}_i)$ term in the loss function ensures that energy value obtained from ground truth and predicted value are as close as possible.

5 RESULTS

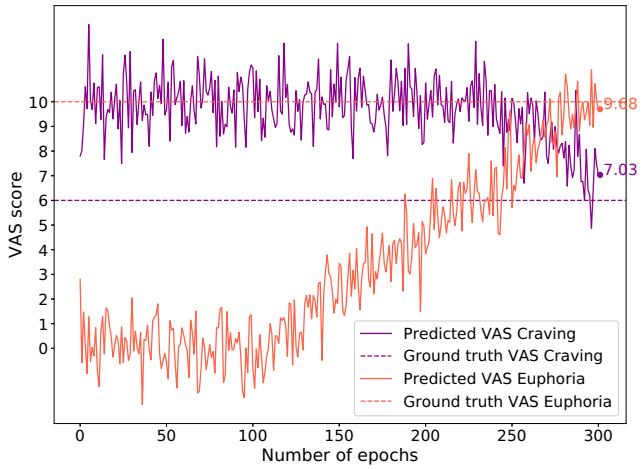

Figure 5: Plot demonstrating how the predicted scores for craving and euphoria are updated during the inference stage for one particular data point. The final predicted scores for craving and euphoria are 7.03 and 9.68

In this section, we review the performance of the SPEN model in predicting VAS craving and euphoria against different baselines. There are mainly two types of model architecture we considered in our experiments:- 1) *Single-label*, where a separate model was used for predicting VAS craving and euphoria. Both these models were trained independent of the other, i.e.; there is no parameter sharing between these models, 2) *Multi-label*, where both VAS craving and euphoria

		VAS Craving			VAS Euphoria		
Model	**Prediction type**	ρ	*NRMSE(%)*	*NMAE(%)*	ρ	*NRMSE*	*NMAE*
Baseline-Linear Regression[11]	Single-label	0.41	18.6	14.6	0.72	19.9	15.8
Linear Regression	Multi-label	0.37	22.9	18.7	0.57	24	19.8
Neuural Network (2 lnn)	Multilabel	0.41	21.8	18.2	0.62	22.5	18.4
MTL	Multi-label	0.42	18.1	14.2	0.72	19.4	15.7
SPEN	Multi-label	**0.51**	**14.2**	**10.4**	**0.78**	**14.5**	**10.8**

Table 1: The performance of VAS Craving and VAS Euphoria trained with different models and architecture types using Leave-One-Subject-Out Cross-Validation (LOSOXV) experiments. The performance was measured in terms of average Pearson correlation coefficient (ρ), Normalized Root-Mean Square Error (NRMSE), Normalized Mean Absolute Error (NMAE) across participants.

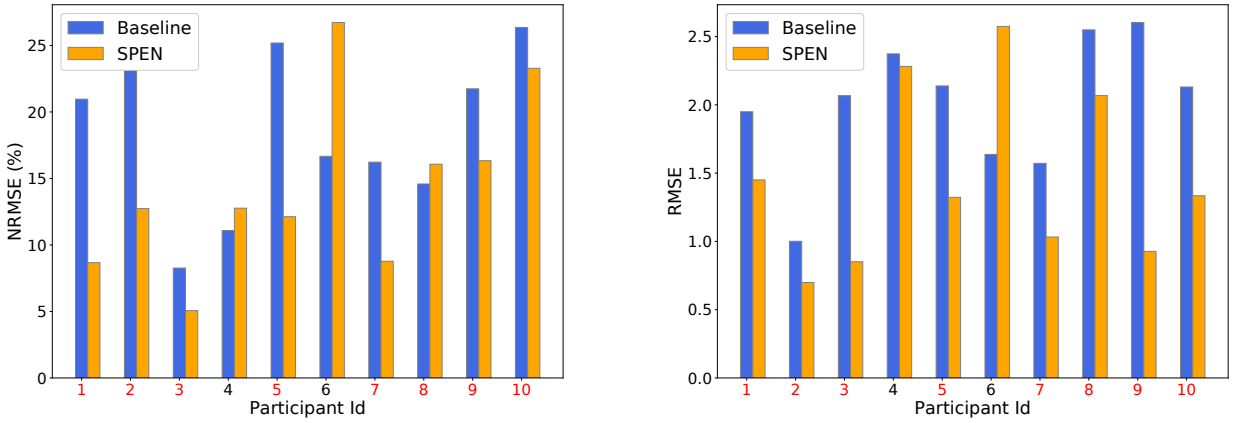

(a) VAS Craving **(b) VAS Euphoria**

Figure 6: shows the performance of our SPEN model against the Baseline for (a) the VAS Craving and (b) the VAS Euphoria across different participants with respect to the NRMSE. For participants 4 and 6, the ground craving and euphoria scores were found not to be correlated with each other. As a result, we could see our SPEN model's performance is worse than the Baseline in these two participants.

are jointly learned using a single model. Table 1 contains the results of all the experiments.

As a baseline, we used the original model proposed by [11], where the authors used a simple linear regression model to predict VAS craving and euphoria separately. We compared this baseline model against models which jointly predict both VAS craving and euphoria. For this purpose, we considered linear regression, 2-layer neural network (lnn), Multi-task learning network (MTL) [6], and finally, SPEN. The linear regression and 2-lnn were implemented straightforward using scikit-learn [21]. The multi-task network for multi-label prediction uses an architecture where each label's prediction is considered as a separate task, and there are a set of global parameters shared between all the tasks and local parameters specific to a task. The input to the model is first passed to a network consisting of global parameters to generate a feature representation. This representation is then passed to models unique to tasks. In all our experiments, we used the same set of features as mentioned in the previous section.

From table 1 we can see that both linear regression and 2-lnn architectures based on multi-label prediction performed worse than the baseline model. The performance of the Multi-task learning-based model was slightly better than the baseline model with just an improvement of $\approx 0.05\%$ in the normalized root-mean-square error (NRMSE). Finally, SPEN architecture that jointly predicts craving and euphoria scores while considering the dependencies between these two variables outperforms baseline and the remaining models. For craving, SPEN architecture results in a NRMSE of 14.2, $\approx 4.5\%$ overall improvement compared to the baseline model of [11] and a Pearson correlation coefficient (ρ) of 0.51, 10% improvement over the baseline model. Similarly, for euphoria,

SPEN architecture achieves NRMSE of 14.5, ≈ 5.5% overall improvement compared to the baseline model and a (ρ) of 0.78, 6% improvement over the baseline model.

As mentioned in the previous section, unlike the traditional machine learning architectures, the output labels are learned during the testing phase in SPENs. To demonstrate how the output labels are updated during inference, in figure 5 we plot predicted craving and euphoria score over the epochs during the inference stage for one data point. For this particular case, the ground truth VAS craving score is 6, and the VAS euphoria score is 10 (highlighted by the horizontal dash lines). The SPEN initially starts with random values for craving and euphoria scores ($\bar{y}$), and then performs gradient descent to update these scores for minimizing the energy. In this case, the final predicted craving and euphoria scores for this data point are 7.03 and 9.68, respectively.

In the next step, we analyzed the performance of the SPEN architecture against the baseline model for all the participants. Figure 6 shows a barplot comparing the performance between these model architectures for VAS craving and euphoria prediction. The participants in which ground-truth VAS craving and euphoria are correlated were highlighted (all participants except for 4 and 6) in red color. From this figure, we can conclude that SPEN architecture outperformed the baseline model in all the participants where the labels are correlated. It performed almost the same or worse than baseline in participants with no label correlation. This shows that SPEN architecture, along with associating an input with the appropriate output variable, could leverage the correlation information present between variables in the dataset to improve both craving and euphoria performance.

6 CONCLUSION AND FUTURE WORK

While the previous work by [11] had shown that cardiac signals could indeed predict cocaine craving and euphoria, we show that we can improve upon this performance if we use a system that considers the relationship between output labels while also considering the cardiac signal. To achieve this, we used SPENs, a deep architecture that defines an energy function on the input and the candidate labels, and the predictions are then obtained by approximately minimizing the energy via gradient descent. Multilabel prediction with label Correlations has been well studied in computer vision and neural language procession, but to the best of our knowledge, its applicability in ubiquitous technologies is limited even though we encounter it in many problem scenarios. Some examples include:- activity recognition and duration of the activity, type and amount of drug administration recognition, human context recognition which includes location, emotion, activity, etc. As a next step, we want to extend SPEN architecture to these cases.

In the previous section, we had seen how SPEN performed poorly on participants whose labels are not correlated. This is a significant limitation of our work, as essentially SPEN fails on participants whose labels do not follow the trend of a general population. One way to overcome this is by grouping subjects based on their label dependencies and building systems unique to each subpopulation. Unfortunately, we could not try this because our current dataset is collected only from a small set of participants, and creating subpopulations means even lesser data is available for the model to learn. Another limitation of our work is that SPEN requires learning of the labels by performing gradient descent even at the inference stage. As a result, its practical applicability is limited. We plan to explore other directions that can overcome these while still making predictions by considering label correlations.

REFERENCES

[1] 2015. Zephyr Bioharness 3. https://www.zephyranywhere.com/media/download/bioharness3-user-manual.pdf.

[2] Paul Baskett, Yi Shang, Michael V Patterson, and Timothy Trull. 2013. Towards a system for body-area sensing and detection of alcohol craving and mood dysregulation. In *2013 IEEE 10th Consumer Communications and Networking Conference (CCNC)*. IEEE, 875–876.

[3] David Belanger and Andrew McCallum. 2016. Structured prediction energy networks. In *International Conference on Machine Learning*. PMLR, 983–992.

[4] Alan S Campbell, Jayoung Kim, and Joseph Wang. 2018. Wearable electrochemical alcohol biosensors. *Current opinion in electrochemistry* 10 (2018), 126–135.

[5] Stephanie Carreiro, Peter R Chai, Jennifer Carey, Jeffrey Lai, David Smelson, and Edward W Boyer. 2018. mHealth for the detection and intervention in adolescent and young adult substance use disorder. *Current addiction reports* 5, 2 (2018), 110–119.

[6] Rich Caruana. 1997. Multitask learning. *Machine learning* 28, 1 (1997), 41–75.

[7] Soujanya Chatterjee, Karen Hovsepian, Hillol Sarker, Nazir Saleheen, Mustafa al'Absi, Gowtham Atluri, Emre Ertin, Cho Lam, Andrine Lemieux, Motohiro Nakajima, et al. 2016. mCrave: continuous estimation of craving during smoking cessation. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. 863–874.

[8] Guibin Chen, Deheng Ye, Zhenchang Xing, Jieshan Chen, and Erik Cambria. 2017. Ensemble application of convolutional and recurrent neural networks for multi-label text categorization. In *2017 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2377–2383.

[9] Zhao-Min Chen, Xiu-Shen Wei, Peng Wang, and Yanwen Guo. 2019. Multi-label image recognition with graph convolutional networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5177–5186.

[10] Charles Dackis and Charles O'Brien. 2005. Neurobiology of addiction: treatment and public policy ramifications. *Nature neuroscience* 8, 11 (2005), 1431.

[11] Bhanu Teja Gullapalli, Annamalai Natarajan, Gustavo A Angarita, Robert T Malison, Deepak Ganesan, and Tauhidur Rahman. 2019. On-body sensing of cocaine craving, euphoria and drug-seeking behavior using cardiac and respiratory signals. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3, 2 (2019),

1–31.

[12] Syed Monowar Hossain, Amin Ahsan Ali, Md Mahbubur Rahman, Emre Ertine David Epstein, Ashley Kennedy, Kenzie Preston, Annie Umbricht, Yixin Chen, and Santosh Kumar. 2014. Identifying drug (cocaine) intake events from acute physiological response in the presence of free-living physical activity. In *IPSN-14 Proceedings of the 13th International Symposium on Information Processing in Sensor Networks*. IEEE, 71–82.

[13] Sheng-Jun Huang and Zhi-Hua Zhou. 2012. Multi-label learning by exploiting label correlations locally. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 26.

[14] Masudul H Imtiaz, Volkan Y Senyurek, Prajakta Belsare, Stephen Tiffany, and Edward Sazonov. 2019. Objective detection of cigarette smoking from physiological sensor signals. In *2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 3563–3566.

[15] Feng Kang, Rong Jin, and Rahul Sukthankar. 2006. Correlated label propagation with application to multi-label learning. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, Vol. 2. IEEE, 1719–1726.

[16] Xiangnan Kong, Bokai Cao, and Philip S Yu. 2013. Multi-label classification by mining label and instance correlations from heterogeneous information networks. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*. 614–622.

[17] Md Shaad Mahmud, Hua Fang, Honggang Wang, Stephanie Carreiro, and Edward Boyer. 2018. Automatic Detection of Opioid Intake Using Wearable Biosensor. In *2018 International Conference on Computing, Networking and Communications (ICNC)*. IEEE, 784–788.

[18] Vinod Nair and Geoffrey E Hinton. 2010. Rectified linear units improve restricted boltzmann machines. In *Icml*.

[19] Annamalai Natarajan, Gustavo Angarita, Edward Gaiser, Robert Malison, Deepak Ganesan, and Benjamin M Marlin. 2016. Domain adaptation methods for improving lab-to-field generalization of cocaine detection using wearable ECG. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. 875–885.

[20] Annamalai Natarajan, Abhinav Parate, Edward Gaiser, Gustavo Angarita, Robert Malison, Benjamin Marlin, and Deepak Ganesan. 2013. Detecting cocaine use with wearable electrocardiogram sensors. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*. ACM, 123–132.

[21] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.

[22] Anjan Panneer Selvam, Sriram Muthukumar, Vikramshankar Kamakoti, and Shalini Prasad. 2016. A wearable biochemical sensor for monitoring alcohol consumption lifestyle through Ethyl glucuronide (EtG) detection in human sweat. *Scientific reports* 6, 1 (2016), 1–11.

[23] Volkan Y Senyurek, Masudul H Imtiaz, Prajakta Belsare, Stephen Tiffany, and Edward Sazonov. 2020. A CNN-LSTM neural network for recognition of puffing in smoking episodes using wearable sensors. *Biomedical Engineering Letters* 10, 2 (2020), 195–203.

[24] Ruiqi Shi, Zhang Chen, Haidong Wang, Peng Sun, Timothy Trull, and Yi Shang. 2015. mAAS–A Mobile Ambulatory Assessment System for Alcohol Craving Studies. In *2015 IEEE 39th Annual Computer Software and Applications Conference*, Vol. 3. IEEE, 282–287.

[25] Atapol Sughondhabirom, Diwakar Jain, Ralitza Gueorguieva, Vladimir Coric, Robert Berman, Wendy J Lynch, David Self, Peter Jatlow, and Robert T Malison. 2005. A paradigm to investigate the self-regulation of cocaine administration in humans. *Psychopharmacology* 180, 3 (2005), 436–446.

[26] Yinyuan Zhang, Ricardo Henao, Zhe Gan, Yitong Li, and Lawrence Carin. 2018. Multi-label learning from medical plain text with convolutional residual models. In *Machine Learning for Healthcare Conference*. PMLR, 280–294.