JULIAN RAMOS

ing.julianr@gmail.com | www.julian-ramos.com | Julian Ramos

Research Statement

Personalized medicine (i.e., precision medicine) [2] is an approach where health treatment is adjusted to an individual's genetic, environmental, and lifestyle factors. The main goal is to avoid ineffective treatments and side effects that delay health improvement [10]. Despite the advantages, genetic based personalized medicine is at a basic research stage, and it is not used yet in clinical practice [13]. Personalization is still done by hand, starting from a one-size-fits-all treatment, and then adjusting for each patient in a lengthy trial and error process. My research develops **mobile health (mHealth) technologies for sensing and intervention**. My main interest is in personalized mHealth: re-imagining interventions as a process that adapts automatically to a patient's health and preferences. I develop methods that use artificial intelligence (AI) to optimize interventions for each patient selecting what, when, and how to intervene informed by sensors and human-feedback, in an interactive process. My goal is to create robust technology capable of dealing with the high variability of real-world deployments. With this technology, I want to address pressing societal issues in public and personal health, with minimal or no barriers to access. In the next section, I will describe my research contributions, starting with early work in stress recognition and interruptibility, which inspired some of my work on personalized mHealth interventions. After that, I will describe my future research and how it has been informed by preliminary results from an mHealth intervention study deployed during the pandemic (May - July 2020).

Selected Research Contributions

STRESS DETECTION WHILE EXERCISING Detecting stressful events is very important to prevent relapse for patients undergoing drug and alcohol rehabilitation. When I worked on this area (2014), stress recognition techniques were limited to function in lab conditions (e.g., while the subject was sitting). In my research, I advanced the field by introducing stress recognition methods that work in close to real-life conditions like when the subject is engaged in physical activity (e.g., walking, running, and cycling). Recognizing stress while exercising is challenging because physiological signals necessary for detection like heart rate, breathing rate, and galvanic skin response, change drastically with physical activity. To overcome this challenge, we created a stress recognition detector [7] specialized for the physical activity of the subject. Our solution consisted of using an activity classifier to recognize the participant's activity and then use a stress classifier built from stress data for that specific activity. This work was the first detecting stress in close to real-life scenarios, resulting in over a 20% increase in the detection rate of stressful events compared to using a single classifier and ignoring the current activity. However, this solution is limited because it requires pre-built activity recognition models for any activity a person could engage in. To overcome this new challenge [23], I created a machine learning approach that, instead of using specific activity recognition classifiers, detects whether the physical activity is of high or low exertion. Then it uses a stress recognition model appropriate for it. This research has gained a lot of attention recently due to advances in deep learning and the ubiquity of wearable devices like the Fitbit and Apple watch. My former research on physical activity recognition informed most of this work [9, 8, 6]

INTERRUPTIBILITY DETECTION Interruptibility detection recognizes states when a computer user can be interrupted to receive information or engage in an activity different from the current one. When I worked on Interruptibility detection, the focus was on detecting interruptibility while working on a desktop. However, the advent of the smartphone opened up the opportunity to detect interruptibility anytime the user is near her phone (a significant amount of time [3]). In this work, we developed different methods for detecting interruptibility, leveraging all these new capabilities. Our first method uses smartphone UI events to detect interruptibility, and we found [19] that delivering notifications during interruptible states had half the cognitive burden of a randomly timed notification. Then, we evaluated a multi-device version [20], and we added activity recognition and found that this reduced by three times the cognitive load burden compared to the UI event plus phone-only interruptibility detection. Interruptibility detection research was motivated to enhance productivity and improve user experience; however, as I will describe in the next section, this technique can be adapted to detect the best context to deliver treatment (i.e., a receptivity state) in an mHealth intervention. This interruptibility work resulted in an honorable mention and multiple highly cited articles in top venues like UBICOMP and PERCOM.

PERSONALIZATION OF TIME OF TREATMENT AND CONTENT Personalization of a health intervention helps to avoid adverse side effects and quickly improve primary outcomes. Despite advances, medical interventions are personalized manually by health professionals. Artificial intelligence (AI) has been used for personalizing the content of interventions with promising results [22]. However, these approaches are under-powered by lacking a way to remind participants of treatment. A mechanism to remind participants of treatment is to detect Receptivity: a state in which a patient is most likely to follow advice[18]. Receptivity is correlated with treatment adherence [12]; however, it has not been shown [17] to improve primary outcomes as an active intervention component. In this project, I explored the simultaneous personalization of timing and content in a sleep intervention through an android App called SleepU. I used contextual bandits (AI) for content personalization and built a receptivity detector to personalize timing. The receptivity detector is an extension of my work in interruptibility detection. This approach was validated in a deployment to 30 participants for four weeks in the Spring of 2019. I found that SleepU improved the participants' sleep duration, time in bed, adherence to the sleep recommendations, and motivation to improve their sleep. This combined approach overcomes previous limitations in AI personalization by using receptivity as a mechanism for proactively reminding of treatment. To my knowledge, this work would be the first to demonstrate that this combined approach results in a statistically significant improvement of several outcomes (e.g., primary health outcome, adherence, and motivation) in an mHealth intervention. This work was supported by the **2017 Digital Health Fellowship** (USD$75k) from the Center for Machine Learning and Health (CMLH) and is currently in submission to UBICOMP 2021.

Following these successful results, I was invited by CMLH to submit a full proposal developing further these personalization methods. This proposal is currently under submission and I describe it in my future research section. Also, as a result of this personalization project, I received the **2019 Microsoft Dissertation Grant**, which allowed me to deploy a more extensive study (n=80). In this study, participants interacted for six weeks with a slightly modified versions of the SleepU app used to test in detail the different personalization effects. This study started March 2nd of 2020, only a week before the pandemic-lockdown in Pennsylvania started. Preliminary results show that there was no significant change in sleep measures due to the intervention. Changes in the way participants interacted with their phones and pandemic-induced attitudes towards sleep explain this null result. I will be publishing these findings as guidelines for the adaptation of mobile health interventions for a pandemic. This is very important for the current and future pandemics; Considering that the conditions favorable for the development of zoonotic diseases that cause pandemics like deforestation and urbanization [5, 14] are still present [4] and not projected to improve, the occurrence of the next pandemic is only a matter of time. In the next section, I will describe my future research and how these preliminary results informed it.

Future Research Agenda

As a faculty member, my research agenda will continue to revolve around mobile health and wellness. I want to explore more aspects of personalization like the language style used to communicate with the patient. I'm also interesting in the interaction of language style with contextual factors and its effect on intervention outcomes.

I'm also interested in continuing my research in sensing emotions like stress, depression, and anxiety from physiological sensor data, smartphone usage, and speech. I believe that due to recent advances in transfer-learning and self-supervised learning, it is not possible to create emotion detectors from very small amounts of labeled as it is typical from most deployments. The next is only a handful of the most concrete research ideas I have, including a research grant that has been short-listed for full proposal submission by CMLH:

Beyond receptivity Unusual events (e.g., the pandemic) affects the way people interact with their electronic devices, rendering useless pre-trained models to detect receptivity. Training new or updating models although possible does not work very well[17] and requires at least a couple of weeks of data [12]. In the pandemic scenario, even the receptivity concept that hinges around discovering opportunistic states may be unfeasible. To overcome this challenge, the AI personalizing timing of treatment may need to try different communication channels (chatbots, SMS, phone call, social networks, email). For this project, I would like to explore multi-device (phone, wearable, desktop) receptivity detection and multi-channel health interventions. The main goal is to discover which device and channel work best for each patient using minimal data or feedback. Another important aspect of this project is the study of the effect of context, based on my own research it would be expected that some devices and channels have higher treatment adherence depending on the time of day, activity and location of the participant.

Language-style The particular language-style used for communicating, has improved intervention outcomes in personalized health interventions [21, 24]. There are many different language-styles (e.g., empathetic, authoritative, supportive); however, figuring out the one that works best for each individual is time-consuming and current approaches rely on the response to questionnaires, which limits adaptation. In this project, I want to develop methods for the automatic adaptation of tone based on health outcomes and context. I envision mobile health interventions that learn over time which language-style works better and how context (e.g., time-of-day, weekend vs. weekday) or intervention specific measures (e.g., motivation) affect health outcomes. For example, for low motivation, empathetic messages could work better, while authoritative messages may be more effective at high motivation levels.

Improving cardiovascular function for cancer survivors There are an estimated 17 million cancer survivors in the United States [16]. Cancer survivors are more likely than healthy adults to exhibit [11] late effects of treatment that negatively impact cardiometabolic health. Recent reviews found that physical activity (PA) decreased all-cause and cancer-specific mortality in individuals diagnosed with breast, colorectal, or prostate cancer [15]. Despite these advantages, cancer survivors' adherence to PA recommendations is low [1]. In this project, the leading hypothesis is that personalizing PA interventions using AI and sensors will improve adherence to PA among cancer survivors. I'm the PI of this research proposal, currently under submission at the Center for Machine Learning and Health (USD$300K). I was invited to submit this proposal after presenting my personalization work. This work is a collaboration with **Carissa Low** and **Robert Kraut**.

I firmly believe, my research experience in data collection and analysis from behavioral studies, applied machine learning to infer mental states, and experience deploying mobile health interventions in the real world, puts me in a unique position to lead and contribute to research efforts both in industry and academia.

References

[1] C. M. Blanchard, K. S. Courneya, and K. Stein. Cancer survivors' adherence to lifestyle behavior recommendations and associations with health-related quality of life: results from the american cancer society's scs-ii. *Journal of Clinical Oncology*, 26(13):2198–2204, 2008.

[2] F. S. Collins and H. Varmus. A new initiative on precision medicine. *New England journal of medicine*, 372(9):793–795, 2015.

[3] A. K. Dey, K. Wac, D. Ferreira, K. Tassini, J.-H. Hong, and J. **Ramos**. Getting closer: an empirical investigation of the proximity of user to their smart phones. In *Proceedings of the 13th international conference on Ubiquitous computing*, pages 163–172, 2011.

[4] Food and A. O. of the United Nations. Global forest resources assessment 2020 – key findings. 2020.

[5] R. Gibb, D. W. Redding, K. Q. Chin, C. A. Donnelly, T. M. Blackburn, T. Newbold, and K. E. Jones. Zoonotic host diversity increases in human-dominated ecosystems. *Nature*, 584(7821):398–402, 2020.

[6] H. Gjoreski, S. Kozina, M. Gams, M. Lustrek, J. A. Álvarez-García, J.-H. Hong, J. **Ramos**, A. K. Dey, M. Bocca, and N. Patwari. Competitive live evaluations of activity-recognition systems. *IEEE Pervasive Computing*, 14(1):70–77, 2015.

[7] J.-H. Hong, J. **Ramos**, and A. K. Dey. Understanding physiological responses to stressors during physical activity. In *Proceedings of the 2012 ACM conference on ubiquitous computing*, pages 270–279, 2012.

[8] J.-H. Hong, J. **Ramos**, and A. K. Dey. Toward personalized activity recognition systems with a semipopulation approach. *IEEE Transactions on Human-Machine Systems*, 46(1):101–112, 2015.

[9] J.-H. Hong, J. **Ramos**, C. Shin, and A. K. Dey. An activity recognition system for ambient assisted living environments. In *International Competition on Evaluating AAL Systems Through Competitive Benchmarking*, pages 148–158. Springer, 2012.

[10] J. L. Jameson and D. L. Longo. Precision medicine—personalized, problematic, and promising. *Obstetrical & gynecological survey*, 70(10):612–614, 2015.

[11] R. B. Kim, A. Phillips, K. Herrick, M. Helou, C. Rafie, M. S. Anscher, R. B. Mikkelsen, and Y. Ning. Physical activity and sedentary behavior of cancer survivors and non-cancer individuals: results from a national survey. *PloS one*, 8(3):e57598, 2013.

[12] F. Künzler, V. Mishra, J.-N. Kramer, D. Kotz, E. Fleisch, and T. Kowatsch. Exploring the state-of-receptivity for mhealth interventions. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 3(4), Dec. 2019.

[13] C. M. Lewis and E. Vassos. Polygenic risk scores: from research tools to clinical instruments. *Genome Medicine*, 12:1–11, 2020.

[14] N. Madhav, B. Oppenheim, M. Gallivan, P. Mulembakani, E. Rubin, and N. Wolfe. Pandemics: risks, impacts, and mitigation. 2017.

[15] A. McTiernan, C. M. Friedenreich, P. T. Katzmarzyk, K. E. Powell, R. Macko, D. Buchner, L. S. Pescatello, B. Bloodgood, B. Tennant, A. Vaux-Bjerke, et al. Physical activity in cancer prevention and survival: A systematic review. *Medicine and science in sports and exercise*, 51(6):1252–1261, 2019.

[16] K. D. Miller, L. Nogueira, A. B. Mariotto, J. H. Rowland, K. R. Yabroff, C. M. Alfano, A. Jemal, J. L. Kramer, and R. L. Siegel. Cancer treatment and survivorship statistics, 2019. *CA: a cancer journal for clinicians*, 69(5):363–385, 2019.

[17] L. G. Morrison, C. Hargood, V. Pejovic, A. W. Geraghty, S. Lloyd, N. Goodman, D. T. Michaelides, A. Weston, M. Musolesi, M. J. Weal, et al. The effect of timing and frequency of push notifications on usage of a smartphone-based stress management intervention: an exploratory trial. *PloS one*, 12(1):e0169162, 2017.

[18] I. Nahum-Shani, S. N. Smith, B. J. Spring, L. M. Collins, K. Witkiewitz, A. Tewari, and S. A. Murphy. Just-in-time adaptive interventions (jitais) in mobile health: key components and design principles for ongoing health behavior support. *Annals of Behavioral Medicine*, 52(6):446–462, 2017.

[19] T. Okoshi, J. **Ramos**, H. Nozaki, J. Nakazawa, A. K. Dey, and H. Tokuda. Attelia: Reducing user's cognitive load due to interruptive notifications on smart phones. In *2015 IEEE International Conference on Pervasive Computing and Communications (PerCom)*, pages 96–104. IEEE, 2015.

[20] T. Okoshi, J. **Ramos**, H. Nozaki, J. Nakazawa, A. K. Dey, and H. Tokuda. Reducing users' perceived mental effort due to interruptive notifications in multi-device mobile environments. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 475–486, 2015.

[21] C. M. Pollard, P. A. Howat, I. S. Pratt, C. J. Boushey, E. J. Delp, and D. A. Kerr. Preferred tone of nutrition text messages for young adults: Focus group testing. *JMIR mHealth uHealth*, 4(1):e1, Jan 2016.

[22] M. Rabbi, M. H. Aung, M. Zhang, and T. Choudhury. MyBehavior: Automatic Personalized Health Feedback from User Behaviors and Preferences using Smartphones. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing - UbiComp '15*, pages 707–718, New York, New York, USA, 2016. ACM Press.

[23] J. **Ramos**, J.-H. Hong, and A. K. Dey. Stress recognition-a step outside the lab. In *International Conference on Physiological Computing Systems*, volume 2, pages 107–118. SCITEPRESS, 2014.

[24] E. Yom-Tov, G. Feraru, M. Kozdoba, S. Mannor, M. Tennenholtz, and I. Hochberg. Encouraging physical activity in patients with diabetes: intervention using a reinforcement learning system. *Journal of medical Internet research*, 19(10):e338, 2017.