Draft, accepted for publication: Proc. 8th International Conference on Affective Computing & Intelligent Interaction (ACII 2019)

A Robot Model of Stress-Induced Compulsive Behavior

Matthew Lewis and Lola Cañamero
Embodied Emotion, Cognition and (Inter-)Action Lab.
Dept. of Computer Science, University of Hertfordshire
Hatfield, Hertfordshire AL10 9AB, UK
M.Lewis4@herts.ac.uk, L.Canamero@herts.ac.uk

***Abstract*—Stress is one of the potential mechanisms underlying compulsive behavior in obsessive-compulsive spectrum disorders. In this paper, we present a robot model and experiments investigating the interactions between internally- and externally-induced stress and compulsive behavior. Our results show properties of the model with potential implications for understanding how stress can result in the generation and maintenance of compulsive behaviors, and how response-prevention interventions can affect compulsive responses under different conditions.**

***Index Terms*—affect modeling, robot model, stress, compulsive behavior, OC-spectrum disorders, motivated action selection, hormonal modulation**

I. Introduction

Obsessive-Compulsive Disorder (OCD) is a disabling mental health disorder characterized by obsessions (recurrent, invasive, often unpleasant thoughts) and compulsions (a strong urge to carry out certain repetitive or ritualized behaviors, such as hand washing or excessive checking). OCD is considered as part of the obsessive-compulsive (OC) spectrum, which also includes conditions such as trichotillomania (pathological hair pulling), body dysmorphic disorder (BDD), and tic disorders such as Tourette's syndrome [1].

Stress is a complex construct [2] hypothesized to be involved in the two main characteristics of OCD, namely obsessions [3] and compulsions [4], to the extent that OC-spectrum disorders are considered among the anxiety disorders [5]. Stress is also hypothesized as a mechanism underlying related dysfunctional repetitive behaviors in animals, such as spontaneous stereotypy, possibly induced by unnatural environmental conditions [6].

One of the main treatments for OCD is Cognitive Behavior Therapy (CBT), in particular Exposure and Response Prevention (ERP) [7]. In ERP, the patient is exposed to a stimulus related to their compulsions, but with the help of a therapist they are encouraged not to execute the compulsive behavior. For example, in the case of compulsive hand washing the patient's hand is made dirty and they are supervised during a period in which they are instructed not to wash their hands. While this therapy has proven effective in some cases, it is also a source of stress for patients exposed to a highly unpleasant stimulus that triggers their condition, while they are prevented from executing potentially comforting responses. This negative emotional response can have a detrimental impact on their wellbeing and lead to rejection of the treatment [8].

New research tools are needed to understand the complex interactions among stress, compulsive behavior and ERP therapy. As a step towards this goal, in this paper, we present a robot model and experiments investigating the interactions between internally- and externally-induced stress and compulsive behavior. Our robot can spontaneously develop compulsive "grooming" behavior after being subject to stressful conditions induced experimentally. As part of the experimental conditions, we have designed different types of response-prevention interventions applied under different conditions, and which affect the robot's stress in different ways, with different "therapeutic" outcomes.

The paper is organized as follows: section II presents the robot model, section III describes the experimental setup, section IV discusses the results in the light of understanding and treating OC-spectrum disorders in humans, and finally section V draws some conclusions from this work.

II. Robot Model

The robot model in this paper builds on the model presented in [9]. We used the small two-wheeled Arduino-based Elisa-3 robot[1] (Fig. 1). The Elisa-3 is equipped with a ring of eight infrared (IR) "distance" sensors that allow it to detect objects and obstacles up to a range of approximately 5cm, and four downwards-pointing IR "ground" sensors that allow it to detect dark or light patches beneath it.

The control architecture, which follows the model in [10], is outlined in Fig. 2 and described in the following subsections.

A. Physiological Variables

We gave the robot three physiological variables: energy, integrity and integument. These variables are controlled homeostatically and are analogous to properties that an animal needs to maintain in order to survive. Each variable ranges from 0 to 1000, with 1000 representing its ideal or target value. The robot dies if either energy or integrity falls to zero. The integument variable is inspired by feathers or fur in animals,

ML is supported by an Early Career Research Fellowship on Robots as Embodied Models of Mental Disorders from the University of Hertfordshire awarded to LC.

[1] http://www.gctronic.com/doc/index.php/Elisa-3

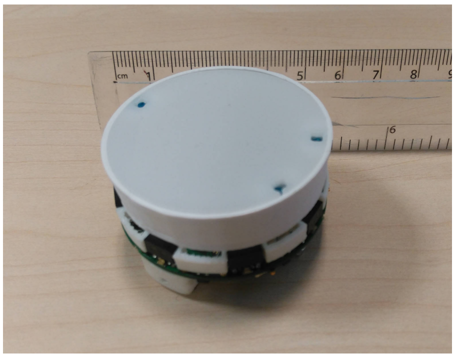

Fig. 1. The Elisa-3 robot.

which if their state fell to a critical state would not directly kill the animal, and for this reason the robot does not die if its integument variable falls to zero. In our model, there can be a mismatch between the "physical" ideal value of a variable and the target perceived by the robot, which is a range of numbers, rather than a single value. The perceived difference between the actual value of a variable and its target range triggers an error signal that will motivate the robot to perform behaviors that correct it. The physiological variables change as follows:

- The energy variable would decay at a constant rate over time, but could be increased by the robot "eating" at "food resources" in the environment (light floor patches).
- The integrity variable would increase at a low constant rate over time as the robot "healed", but damage to the robot was simulated by using the IR distance sensors to detect "contact" and decreasing the integrity variable. When integrity was low, the robot would be more likely to avoid objects in order to avoid further damage.
- The integument variable would decrease slowly over time, but could be increased by the robot "grooming" (causing a set of its IR distance sensors to be activated in sequence, corresponding to an object moving over the sensors) at a "grooming post". In order that the walls of the environment were not available to grooming (making management of this variable an easy task due to the high availability of the required environmental resource), the grooming posts were distinguished from the walls by being positioned on dark patches on the floor (Fig. 3).

B. Hormones

Corresponding to each physiological variable is a simulated hormone that signals the deficit (or error) of that variable, as in [11]. The hormone is released at a rate proportional to the current deficit (the difference between the current value of the variable and the robot's target for that variable) and decays exponentially at a rate proportional to the current level of the hormone. Note that, as we will see below, the target for each physiological variable is not a single value, but rather a range of values within which the deficit is defined as zero. While these simulated hormones are inspired by biological hormones they are not intended to correspond to specific biological hormones, and can be thought of as corresponding to any signaling mechanism that has hormone-like dynamics.

A fourth hormone, released as a function of "stress"—we will call it the "stress hormone"—is released depending on two sources of stress, one internal and the other external. First, stress hormone is released from an internal source in proportion to the level of the three error-signaling hormones. This is similar to the stress mechanism in [11]. Second, stress hormone is released due to an external source when the robot is physically confined. This external source allows us to induce stress in controlled experimental conditions. We took physical enclosure as our source of external stress inspired by the spontaneous development of repetitive stereotypical behavior in animals when kept in stressful conditions [6]. The robot can detect confinement when any pair of its IR distance sensors on opposite sides of its body both detected a nearby object. The stress hormone decays at a slower rate than the error-signaling hormones, so its effects are longer lasting compared to the other hormones which decay more quickly.

The stress hormone acts on the action selection mechanism by changing the target "value" (range) of the physiological variables: low stress expands the range, while high stress contracts it. This mechanism can be thought of as analogous to making the "stressed" robot less tolerant of ambiguity or uncertainly, which has long been theorized to have a link with stress [12]. In our previous work [9], a single target value was used by the robot for each physiological variable, and the difference between the target value and the current value was the deficit. Here, the "target value" is expanded into a "target range" that is modulated by the stress hormone. More precisely, the range is given by

$$R_i = \begin{cases} [t_i - 250 + H_s/4,\ t_i + 250 - H_s/4] & \text{if } H_s \leq 1000 \\ [t_i,\ t_i] & \text{if } H_s > 1000 \end{cases} \tag{1}$$

where i is the relevant error-signaling hormone, t_i is the center of the range (corresponding to the single target value used previously) and H_s is the current value of the stress hormone. The constant values in the above equation were determined empirically prior to the experiments in order that the typical levels for the stress hormone resulted in a target range similar to the values used in [9].

C. Motivations

Three competing motivations, linked to the three physiological variables, guide the behavior of the robot. These motivations are urges to action determined by a combination of the three corresponding internal homeostatically-controlled physiological variables—which provide the robot with "needs"—and by the robot's perception of the environmental resources that can be used to manage those variables.

The values of the three motivations are given by:

$$M_i = H_i + H_i \times \alpha \times C_i \tag{2}$$

where i is the motivation/corresponding hormone, H_i is the current level of hormone i, C_i is the detected size of the

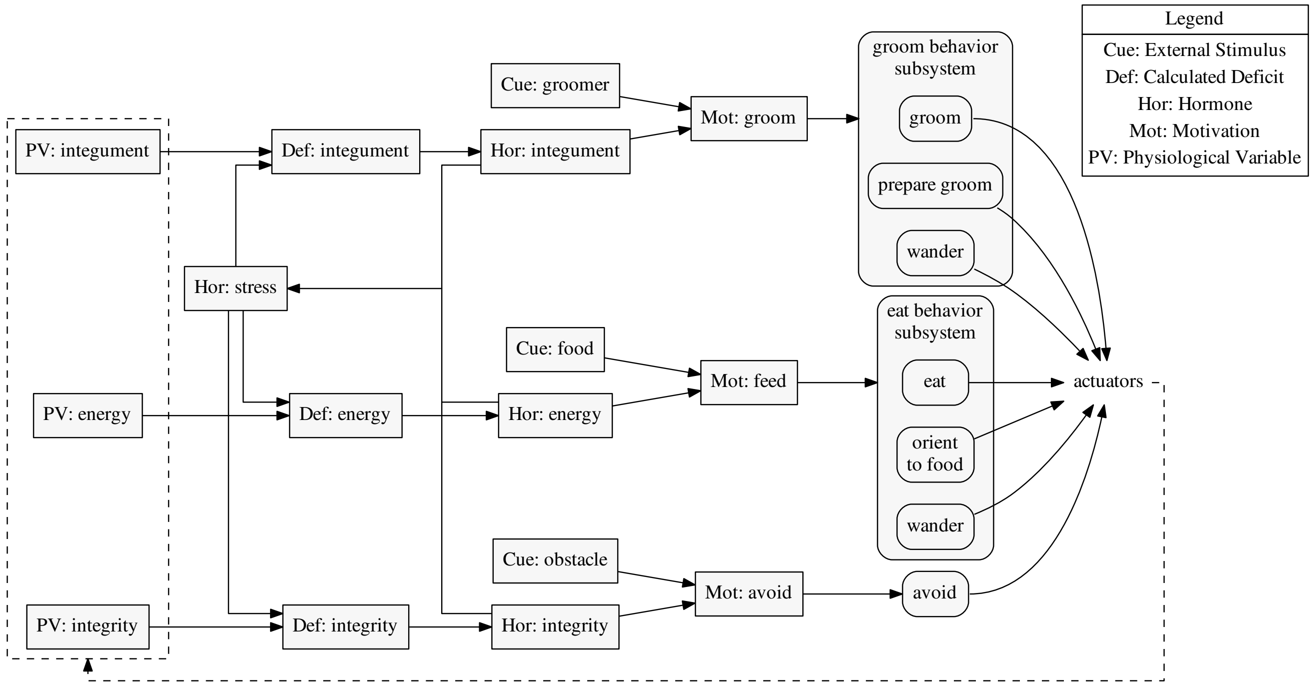

Fig. 2. An outline of the robot's software control architecture. Predefined behaviors (potentially nested) are shown with rounded corners, while other elements of the internal architecture have square corners.

corresponding external cue, and α (here $= 0.05$) is a scaling factor to scale the size of the exteroceptive component.

Equation (2) is similar to that used in our previous work on motivated architectures (e.g. [9], [13]–[15]) inspired by ethology [16, p. 139]. It has been modified here to use a hormonal system for signaling the error, rather than perceiving it directly. The hormonal system means that there is increased hysteresis, however, since these signaling hormones are released and decay rapidly they respond quickly to changes in the deficit.

D. Action Selection

The robot is endowed with a predefined repertoire of behaviors designed to allow it to manage its physiological variables and an action selection algorithm to prioritize them based on the intensity of its motivations.

The action selection algorithm works as follows: At each tick of the action selection mechanism (10Hz), the values indicating the intensity of its three motivations are calculated using (2) and then sorted highest to lowest. Each motivation has an associated behavioral subsystem that includes appetitive (goal-seeking) and consummatory (goal-achieving) sub-behaviors associated with satisfying the corresponding motivation (rounded boxes in Fig. 2). The top-level behavior corresponding to the highest motivation is executed, and it executes any of its sub-behaviors that are executable (behaviors are executable if they are both relevant – "eat" behavior is only relevant if food is detected, otherwise an appetitive behavior to locate food will be relevant – and do not conflict with any behaviors that are already being executed in that action selection tick). Once the top-level behavior with the highest motivation has started all its executable sub-behaviors, the top-level behavior with the next-highest motivation will be executed, and so on.

More than one behavior can be executed simultaneously during an action selection tick. However, behaviors can require one or more actuators in order to execute (the "wander" behavior uses the "wheels" actuator, while the "eat" behavior uses the "mouth" actuator) and behaviors cannot be executed simultaneously if they use the same actuator, meaning that only the higher-priority behavior will be executed in that action-selection tick.

E. Modeling OC-spectrum disorders

In this model, the robot has the potential to show compulsive behaviors as a result of misperceiving the target value of any of the physiological variables "unrealistically" (higher than what is physically achievable) under conditions promoted by high stress. This differs from our model in [9], in which we pre-set unrealistic target values for the integument variable. Unrealistic perception will affect the decision making (action selection) process in the calculation of the deficit leading to higher release of the error-signaling hormone (Fig. 2) and hence increase the intensity of the corresponding motivation (2), leading to pathological compulsive behavior.

III. Robot Experiments

A. Experimental Setup

We tested our model in a decision-making (action selection) task in which the robot has to satisfy its three needs in a

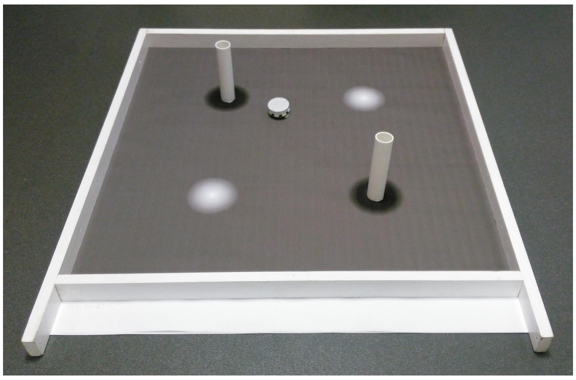

Fig. 3. Our experimental environment. Two grooming posts are positioned on black patches; the two light patches are food resources that provide energy.

timely fashion in order to survive. The robot is placed in an environment containing resources that allow it to satisfy its needs (Fig. 3). This environment consists of an 80cm × 80cm square surrounded by 45mm high wooden walls that can be detected with the robot's distance sensors. The floor is covered in paper, which is printed gray, except for two light areas and two dark areas, which can be detected by the robot's ground sensors, and that indicates the presence of food and grooming resources, respectively. Two 35mm diameter white plastic pipes are fixed in place in the centers of the dark areas to be used as grooming posts.

The internal parameters of the architecture of the robot were set up to have a potential to show compulsive grooming behavior: while the middle value of the target range (t_i in (1)) for its energy and integrity was set 1000 (the "real" ideal value), the middle value of the target range for its integument variable was set to 1150. This value is impossible to achieve, since it is greater than 1000, but previous work with an earlier version of this architecture [9] has suggested that a target value that is beyond what is achievable may be beneficial to the robot as it will make the associated behaviors persist for longer, thus making fuller use of the relevant resource.

The initial values of the physiological variables were: energy: 700; integrity: 900; integument: 700. The initial levels of the hormones were: energy: 500; integrity:100; integument: 500; stress: 500. The initial values of the energy and integument variables were chosen to be high enough that the robot would be unlikely to die shortly after starting due to being unable to find food resources (after finding the first resources, it could then maintain these variables itself). The initial value of the integrity was higher since the robot typically did not damage itself through collisions, so this value was more typical. The initial levels of the hormones were chosen to be approximately equal to the stationary values given the initial values of the physiological variables.

With these initial values, from (1) the robot's target range will be $[1025, \ldots]$, corresponding approximately to the "realistic target" in [9]. Higher levels of the stress hormone will increase the target to the "mildly unrealistic target" value and towards the "highly unrealistic target". We kept the value below the previously used "highly unrealistic target" since that almost always resulted in the robot dying within 6 minutes, and we wanted to capture data over a longer period.

B. Conditions

The same robot was tested in six different conditions, for one run in each condition. The first two conditions were designed as baseline or control conditions for the other four conditions, as follows. In condition 1, there was no externally-induced stress; in condition 2, stress was externally induced, but there was no remedial intervention from the experimenter; in conditions 3–6, the experimenter applied different types of intervention akin to response prevention therapy, as detailed below. In conditions 2–6, stress was externally induced by the experimenter by placing a small plastic container over the robot when it first stopped to feed (Fig. 4). This meant that the robot could survive (due to the food resource) but stress would increase rapidly due its confinement. The plastic container was removed when the level of the stress hormone rose above 1200. Each run lasted 20 minutes, or until the robot died. The precise procedure for each condition was as follows:

- C1. The robot was left in the arena with no interference from the experimenter.
- C2. The robot was "stressed" (see above) shortly after the run started, and then left with no further interference from the experimenter.
- C3. The robot was "stressed" shortly after the run started. Following this, if it groomed and its integument variable increased to 1000 (the maximum value, where the physical need has been totally satisfied) then the experimenter moved it away from the grooming post.
- C4. The robot was "stressed" shortly after the run started. Following this, if it groomed and its energy variable fell below 300 (putting it in danger of dying from neglecting the survival-related need) then the experimenter moved it away from the grooming post.
- C5. The robot was "stressed" shortly after the run started. Following this, if it started grooming then the experimenter moved it away from the grooming post, systematically preventing the behavior linked with compulsions.
- C6. The robot was "stressed" shortly after the run started. Following this, if it groomed and its integument variable increased above 800 then the experimenter moved it away from the grooming post. Here, the physical need has been highly, but not totally, satisfied, unlike in C3.

In conditions 3 to 6, the experimenter moved the robot away from grooming posts by picking it up and quickly moving it about 5cm away from the post.

C. Data Collection

During each run the radio connection was used to transmit to a PC the values of the physiological variables, the levels of the four hormones, the levels of the three motivations, the wheel speeds, and the currently active behaviors. These values were recorded every 250ms.

Fig. 4. The Elisa-3 robot covered with a plastic container, to induce stress in conditions 2–6. The robot was trapped with a food resource (light floor patch) so that it did not die through lack of energy while trapped.

IV. Results and Discussion

Fig. 5 shows the values of the three physiological variables and the stress hormone over each run. Let us highlight the main results for each condition.

C1 shows that, if left to interact in the environment, the robot can survive the full 20 minutes. In the absence of externally-induced stress, and with the initial values of stress, the level of the stress hormone has a tendency to self-maintain as a result of the dynamics of the interaction of the robot as it behaves to correct its deficits, i.e. the initial state is somewhat stable. Near the end of the run, after 1000s, a fall in both the energy and integument variables leads to an increase in stress, illustrating that the stress hormone responds as expected.

C2 shows that, with the addition of externally induced stress, the level of stress hormone can reach a high state that also has a tendency to self-maintain after the stressor has disappeared (rather than falling back to the stable low stress level from C1). In this case, the high stress led to pathological levels of grooming, which caused the robot to die at 17m27s as it neglected its survival-related need for energy.

C3 shows that stopping over-grooming (but still allowing the robot to groom until reaching the maximum value of the integument variable) resulted in the stress falling over time. After 850s, the robot would stop grooming by itself when the integument reached its maximum value, and the experimenter did not need to take action. This can be viewed as a successful intervention of our response prevention therapy. However, we can expect that a repeated stressful event would again result in compulsive behavior, as it might do in animals or humans.

C4 shows that stopping grooming when over-grooming had put the robot in danger of dying—as its energy fell to dangerous values—allowed the robot to survive (unlike C2, where the robot died) but kept the levels of the stress hormone high; therefore, the robot would still show compulsive behavior if the experimenter did not prevent it. This intervention allowed the robot to survive, but did not eliminate compulsive behavior.

C5 represents an extreme version of "response prevention": the robot was stopped from grooming as soon as it started, and hence its integument variable remained near-zero throughout the run. While the robot maintained its energy variable at a high value, and thus survived, its stress variable increased up to an extremely high level of 1804 by the end of the 20 minutes, so it would still groom compulsively if not prevented from doing so, and it would take longer for the level of the stress hormone to fall back to a low value. This intervention could be viewed as counter-productive, since it made a spontaneous return to non-pathological grooming more difficult.

C6 is a variation on C3, with grooming stopped before the integument reached its ideal value of 1000. In this case, the level of the stress hormone fell slower than in C3, indicating that it took longer to bring the robot back to its non-pathological state, but during this time it maintained its variables in a more balanced way—its energy never fell below 400, and all three physiological variables had similar values. By the end of the run, the level of the stress hormone had fallen below 500, so the robot was back to its initial state. This response-prevention intervention was more successful than in C3, since it also allowed the robot to maintain a better "overall health", as its stress level fell to its initial low value.

These experimental results illustrate some properties of our robot model with potential implications for understanding how positive feedback can result in the maintenance of compulsive behaviors, and how response-prevention interventions can affect compulsive responses under different conditions. We have already commented on the latter under C3–C6. Turning to the properties, the most significant one is the self-sustaining dynamics of stress in both low-stress and high-stress conditions. The self-maintaining aspect of high stress, and the corresponding pathological grooming behavior, may be due to a positive feedback loop in which the high stress leads to a small target range for the physiological variables, which makes the perceived deficits larger, leading to release of the error-signaling hormones, and hence of the stress hormone.

V. Conclusion

In this paper, we have presented a robot model and experiments investigating the influence of stress on compulsive behavior, characteristic of OC-spectrum disorders. Our experimental results have illustrated a number of properties of our robot model with potential implications for understanding how positive feedback can result in the maintenance of compulsive behaviors, and how response-prevention interventions can affect compulsive responses under different conditions. One of these properties is the self-sustaining dynamics of stress in both low-stress and high-stress conditions. Regarding the response-prevention interventions, our results show that they can have both beneficial and noxious effects. Some types of response prevention interventions (C3 and C6, stopping grooming when the integument was high) were beneficial as they led to the level of the robot's stress hormone falling back to a level at which the grooming behavior was no longer pathological. Other types of response-prevention interventions were counterproductive as they either resulted in the level of the robot's stress hormone either remaining high (C4, stopping

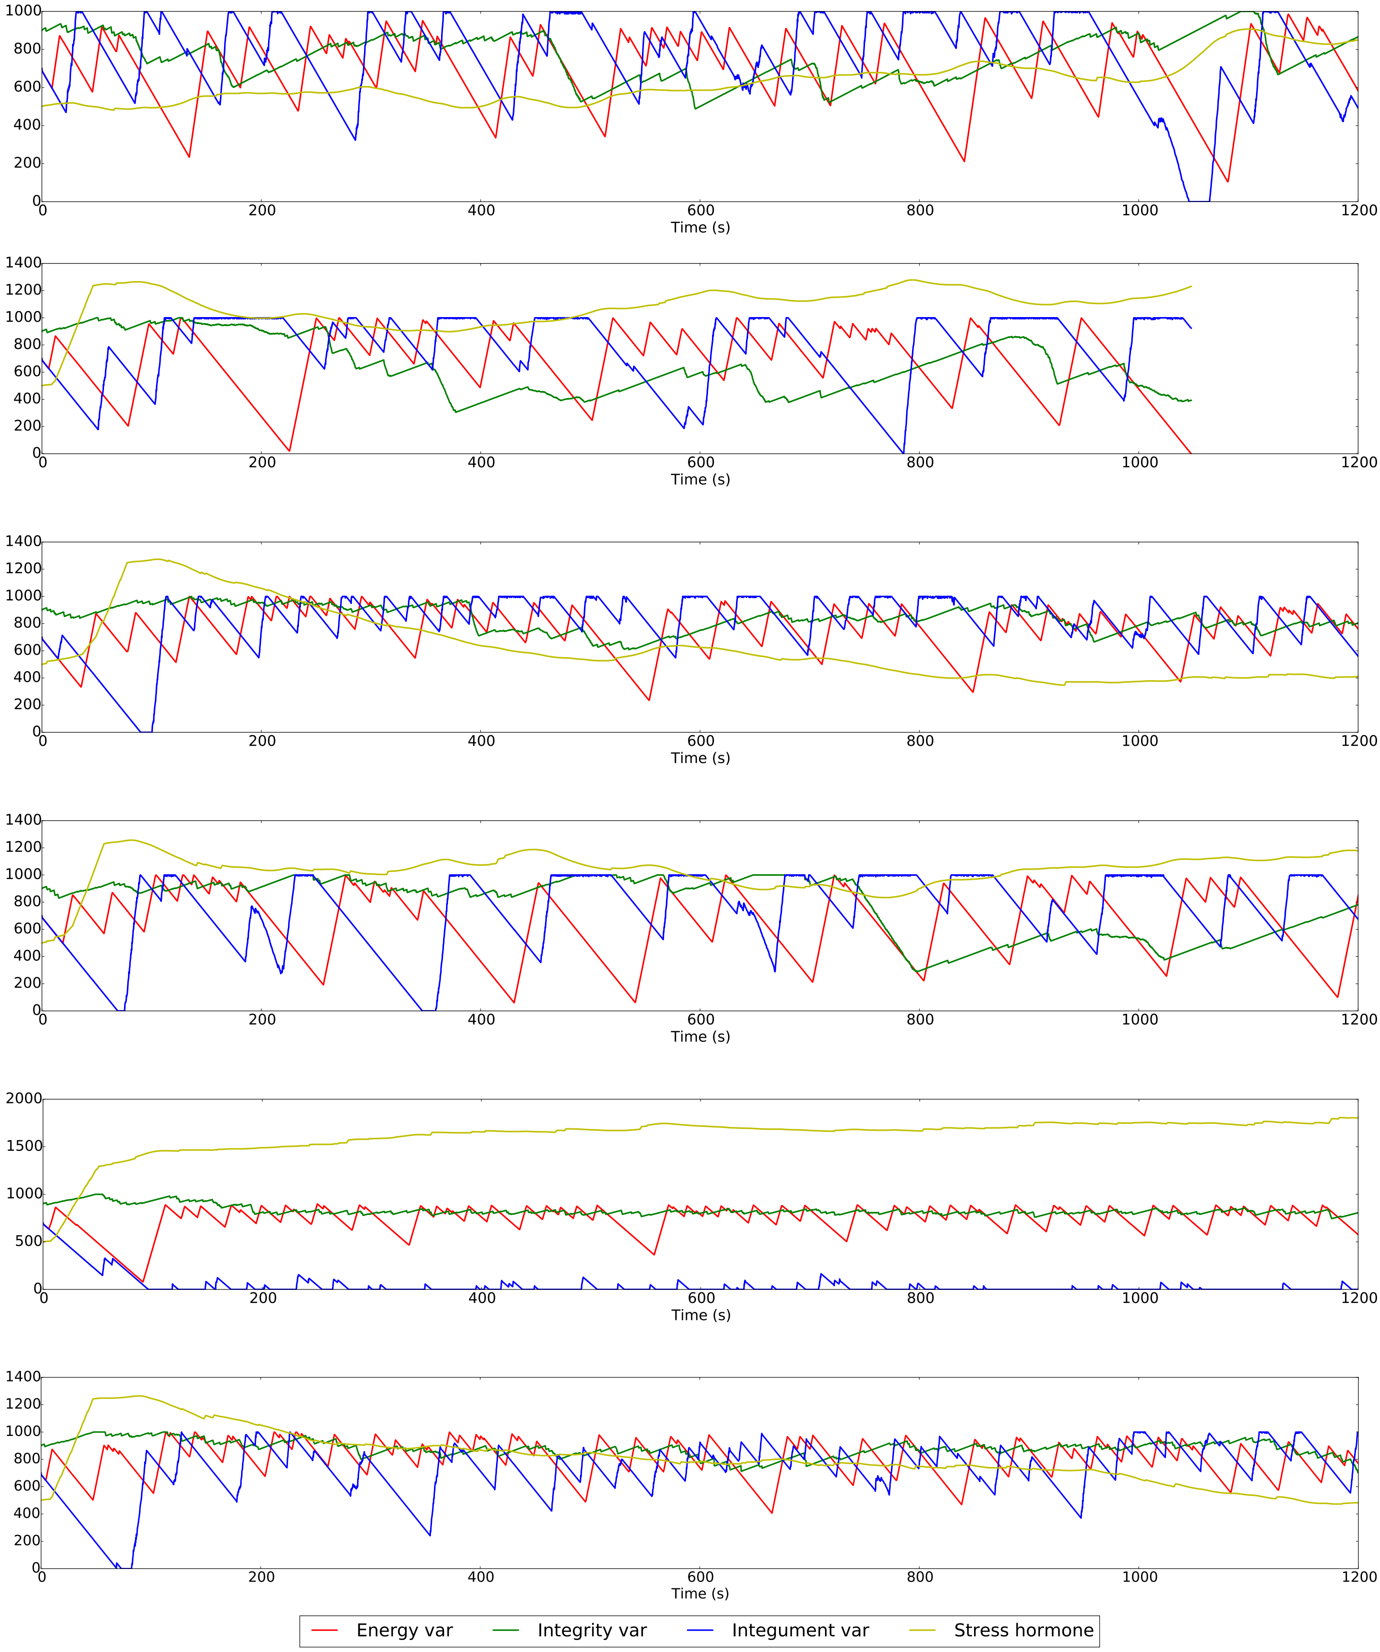

Fig. 5. Values of the three physiological variables and the level of the stress hormone during each of the six runs (C1–C6, top to bottom). In C2, the robot died at 17m27s since its energy fell to zero. In C2–C6, the level of the stress hormone rises sharply near the beginning of the run, when the robot was trapped in the box.

grooming when the energy level was low) or increasing even further (C5, preventing the robot from grooming at all). However, there are differences between our ERP-like interventions and actual ERP therapy—for example, ERP patients are asked to consciously resist performing the ritual behavior, while our robot does not have the cognitive capability to do this and instead the experimenter physically prevents it from grooming.

While these results are preliminary, our model can potentially shed light on the dynamics of interaction between stress and compulsions, and provide a basis to generate and test hypotheses about response-prevention interventions in relevant human and animal pathologies. Future work, with our collaborators in mental health research and treatment, will go in this direction.

Acknowledgment

We would like to thank our collaborators, in particular Prof. Naomi Fineberg (Herts. Partnership NHS Foundation Trust University of Hertfordshire, and Cambridge University).

References

[1] American Psychiatric Association, *Diagnostic and statistical manual of mental disorders: DSM-5*, 5th ed. Arlington, VA: American Psychiatric Association, 2013.

[2] H. Selye, "The evolution of the stress concept," *American Scientist*, 1973.

[3] D. F. Tolin, P. Worhunsky, and N. Maltby, "Are "obsessive" beliefs specific to OCD?: A comparison across anxiety disorders," *Behaviour Research and Therapy*, vol. 44, no. 4, pp. 469–480, 2006. [Online]. Available: https://www.sciencedirect.com/science/article/abs/pii/S0005796705000677

[4] N. A. Fineberg, S. R. Chamberlain, E. Hollander, V. Boulougouris, and T. W. Robbins, "Translational approaches to obsessive-compulsive disorder: From animal models to clinical treatment," *British Journal of Pharmacology*, vol. 164, no. 4, pp. 1044–1061, 2011. [Online]. Available: https://bpspubs.onlinelibrary.wiley.com/doi/full/10.1111/j.1476-5381.2011.01422.x

[5] D. J. Stein, N. A. Fineberg, O. J. Bienvenu, D. Denys, C. Lochner, G. Nestadt, J. F. Leckman, S. L. Rauch, and K. A. Phillips, "Should OCD be classified as an anxiety disorder in DSM-V?" *Depression and Anxiety*, vol. 27, no. 6, pp. 495–506, 2010. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/da.20699

[6] S. D. McBride and M. O. Parker, "The disrupted basal ganglia and behavioural control: an integrative cross-domain perspective of spontaneous stereotypy," *Behavioural brain research*, vol. 276, pp. 45–58, 2015. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S016643281400360X

[7] E. B. Foa, "Cognitive behavioral therapy of obsessive-compulsive disorder," *Dialogues in Clinical Neuroscience*, vol. 12, no. 2, pp. 199–207, 2010. [Online]. Available: https://www.dialogues-cns.org/wp-content/uploads/issues/12/DialoguesClinNeurosci-12-199.pdf

[8] J. S. Abramowitz, "The psychological treatment of obsessive-compulsive disorder," *The Canadian Journal of Psychiatry*, vol. 51, no. 7, pp. 407–416, 2006. [Online]. Available: https://journals.sagepub.com/doi/10.1177/070674370605100702

[9] M. Lewis, N. Fineberg, and L. Cañamero, "A robot model of OC-spectrum disorders: Design framework, implementation and first experiments," *Computational Psychiatry*, 2019. [Online]. Available: https://www.mitpressjournals.org/doi/abs/10.1162/cpsy_a_00025

[10] L. D. Cañamero, "Modeling motivations and emotions as a basis for intelligent behavior," in *Proceedings of the First International Symposium on Autonomous Agents (Agents'97)*, W. L. Johnson, Ed., The ACM Press. Marina del Rey, CA, USA: The ACM Press, 1997, pp. 148–155.

[11] J. Lones, M. Lewis, and L. Cañamero, "A hormone-driven epigenetic mechanism for adaptation in autonomous robots," *IEEE Transactions on Cognitive and Developmental Systems*, vol. 10, pp. 445–454, 2018. [Online]. Available: https://ieeexplore.ieee.org/document/8115310

[12] C. D. Smock, "The influence of psychological stress on the "intolerance of ambiguity"," *The Journal of Abnormal and Social Psychology*, vol. 50, no. 2, pp. 177–182, 1955. [Online]. Available: https://psycnet.apa.org/record/1956-00627-001

[13] L. Cañamero and O. Avila-García, "A bottom-up investigation of emotional modulation in competitive scenarios," in *Proc. Second International Conference on Affective Computing and Intelligent Interaction (ACII 2007)*, ser. Lecture Notes in Computer Science, A. C. R. Paiva, R. Prada, and R. W. Picard, Eds., vol. 4738, Springer Berlin Heidelberg. Lisbon, Portugal: Springer Berlin Heidelberg, 2007, pp. 398–409. [Online]. Available: https://link.springer.com/chapter/10.1007%2F978-3-540-74889-2_35

[14] M. Lewis and L. Cañamero, "An affective autonomous robot toddler to support the development of self-efficacy in diabetic children," in *Proc. 23rd Annual IEEE International Symposium on Robot and Human Interactive Communication (IEEE RO-MAN 2014)*, IEEE. Edinburgh: IEEE, 2014, pp. 359–364. [Online]. Available: https://ieeexplore.ieee.org/document/6926279

[15] ——, "Hedonic quality or reward? A study of basic pleasure in homeostasis and decision making of a motivated autonomous robot," *Adaptive Behavior*, vol. 24, pp. 267–291, 2016. [Online]. Available: https://journals.sagepub.com/doi/full/10.1177/1059712316666331

[16] T. Tyrrell, "Computational mechanisms for action selection," Doctoral dissertation, University of Edinburgh, 1993.