

# A Benchmark Environment Motivated by Industrial Control Problems

Daniel Hein<sup>\*†§</sup>, Stefan Depeweg<sup>\*†§</sup>, Michel Tokic<sup>†</sup>, Steffen Udluft<sup>†</sup>,  
Alexander Hentschel<sup>‡</sup>, Thomas A. Runkler<sup>\*†</sup> and Volkmar Sterzing<sup>†</sup>

<sup>\*</sup>Technische Universität München, Department of Informatics, Boltzmannstr. 3, 85748 Garching, Germany

<sup>†</sup>Siemens AG, Corporate Technology, Otto-Hahn-Ring 6, 81739 Munich, Germany

<sup>‡</sup>AxiomZen, 980 Howe St #350, Vancouver, BC V6Z 1N9, Canada

**Abstract**—In the research area of reinforcement learning (RL), frequently novel and promising methods are developed and introduced to the RL community. However, although many researchers are keen to apply their methods on real-world problems, implementing such methods in real industry environments often is a frustrating and tedious process. Generally, academic research groups have only limited access to real industrial data and applications. For this reason, new methods are usually developed, evaluated and compared by using artificial software benchmarks. On one hand, these benchmarks are designed to provide interpretable RL training scenarios and detailed insight into the learning process of the method on hand. On the other hand, they usually do not share much similarity with industrial real-world applications. For this reason we used our industry experience to design a benchmark which bridges the gap between freely available, documented, and motivated artificial benchmarks and properties of real industrial problems. The resulting industrial benchmark (IB) has been made publicly available to the RL community by publishing its Java and Python code, including an OpenAI Gym wrapper, on Github. In this paper we motivate and describe in detail the IB's dynamics and identify prototypic experimental settings that capture common situations in real-world industry control problems.

## I. INTRODUCTION

Applying reinforcement learning (RL) methods to industrial systems, such as in process industry like steel processing [1], pulp and paper processing [2], and car manufacturing [3], or power generation with gas or wind turbines [4], [5], is an exciting area of research. The hope is that an intelligent agent will provide greater energy efficiency and, desirably, less polluting emissions. However, the learning process also entails a significant amount of risk: we do not know beforehand how a particular learning algorithm will behave, and with complex and expensive systems like these, experiments can be costly.

Therefore, there is high demand in having simulations that share some of the properties that can be observed in these industrial systems.

The existing simulation benchmarks have lead to great advancements in the field of RL. Traditionally simple dynamical systems like pendulum dynamics are studied, whereas nowadays the focus has shifted towards more complex simulators, such as video game environments [6]. Also in the field of robotics very sophisticated simulation environments exist, on which new learning algorithms can be tested [7], [8]. The existence of such benchmarks has played a vital role in pushing the frontier in this domain of science.

For industrial control, however, such a test bed is lacking. In these systems we observe a combination of properties that usually are not present in existing benchmarks, such as high dimensionality combined with complex heteroscedastic stochastic behavior. Furthermore, in industrial control different experimental settings are of relevance, for instance, the focus is usually less on exploration, and more on batch RL settings [9].

To this end, we recently developed the *industrial benchmark* (IB), an open source software benchmark<sup>1</sup>, with both Java and Python implementations, including an OpenAI Gym wrapper, available. Previously, this benchmark has already been used to demonstrate the performance of a particle swarm based RL policy approach [10]. The contribution of this paper lies in presenting the complete benchmark framework as well as mathematical details, accompanied by illustrations and motivations for several design decisions. The IB aims at being realistic in the sense that it includes a variety of aspects that we found to be vital in industrial applications like optimization and control of gas and wind turbines. It is not designed to be an approximation of any real system, but

<sup>§</sup> Equal contributions.

<sup>1</sup>Java/Python source code: <http://github.com/siemens/industrialbenchmark>

to pose the same hardness and complexity. Nevertheless, the process of searching for an optimal action policy on the IB is supposed to resemble the task of finding optimal valve settings for gas turbines or optimal pitch angles and rotor speeds for wind turbines.

The state and action spaces of the IB are continuous and high-dimensional, with a large part of the state being latent to the observer. The dynamical behavior includes heteroscedastic noise and a delayed reward signal that is composed of multiple objectives. The IB is designed such that the optimal policy will not approach a fixed operation point in the three steerings. All of these design choices were driven by our experience with industrial challenges.

This paper has three key contributions: in Section II, we will embed the IB in the landscape of existing benchmarks and show that it possesses a combination of properties other benchmarks do not provide, which makes it a useful addition as a test bed for RL. In Section III we will give a detailed description of the dynamics of the benchmark. Our third contribution, described in Section IV, is to define prototype experimental setups that we find relevant for industrial control. Our goal is to encourage other researchers to study scenarios common in real-world situations.

## II. PLACEMENT OF THE INDUSTRIAL BENCHMARK IN THE RL BENCHMARK DOMAIN

In the RL community numerous benchmark suits exist, on which novel algorithms can be evaluated. For research in industrial control we are interested in a particular set of properties, such as stochastic dynamics with high dimensional continuous state and action spaces. We argue that only few freely available benchmarks exist fulfilling these properties, thereby making our contribution, the IB, a useful addition. To that end, we want to briefly review existing benchmarks in RL.

Classic control problems in RL literature [11], such as the cart-pole balancing and mountain car problems, usually have low dimensional state and action spaces and deterministic dynamics. In the field of robotics more complex and high-dimensional environments exist with a focus on robot locomotion, such as the MuJoCo environment [7], [8]. Other examples are helicopter flight<sup>2</sup> [12] or learning to ride a bicycle [13]. These systems, while complex, usually have deterministic dynamics or only limited observation noise.

Utilizing games as RL benchmarks recently brought promising results of deep RL into the focus of a broad audience. Famous examples include learning to play Atari games<sup>3</sup> [6], [14] based on raw pixels, achieving above-human performance playing Ms. Pac-Man [15], and beating human experts in the game of Go [16]. In these examples, however, the action space is discrete and insights from learning to play a game may not translate to learning to control an industrial system like gas or wind turbines.

<sup>2</sup><https://sites.google.com/site/rlcompetition2014/domains/helicopter>

<sup>3</sup><https://github.com/mgbellemare/Arcade-Learning-Environment>

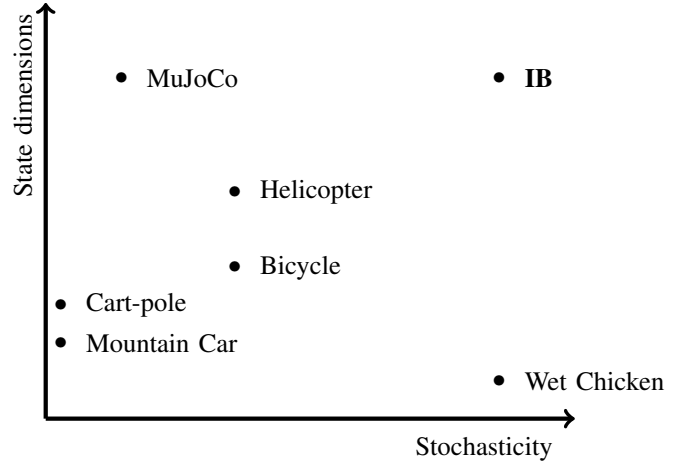


Fig. 1. Qualitative comparison of different RL benchmarks with continuous actions. The state space of the wet chicken 2D benchmark [17] is rather low, but it is highly stochastic which makes it a challenging RL problem. Cart-pole and mountain car are deterministic benchmarks with few state dimensions and only a single action variable. The bicycle benchmark introduces some noise to simulate imperfect balance. The helicopter software simulation has a 12-dimensional state space and a 4-dimensional continuous action space. Stochasticity is introduced to simulate wind effects on the helicopter. The state space of the IB is high, since multiple past observations have to be taken into account to approximate the true underlying Markov state. Stochasticity is not only introduced by adding noise on different observations, but also by stochastic state transitions on hidden variables.

In Figure 1 we give a qualitative overview on the placement of the proposed IB with respect to other RL benchmarks for continuous control. Here, we focus on stochasticity and dimensionality of the benchmarks at hand. Note that by stochasticity we do not only refer to the signal to noise ratio, but also to the structural complexity of the noise, such as heteroscedasticity or multimodality. We conclude that the IB is a useful addition to the set of existing RL benchmarks. In particular, Figure 1 illustrates that the combination of high dimensionality and complex stochasticity appears to be novel compared to existing environments. In the following section, a detailed description and motivation for the applied IB dynamics is presented.

## III. DETAILED DESCRIPTION

At any time step  $t$  the RL agent can influence the environment, i.e., the IB, via actions  $a_t$  that are three dimensional vectors in  $[-1, 1]^3$ . Each action can be interpreted as three proposed changes to the three observable state variables called current steerings. Those current steerings are named velocity  $v$ , gain  $g$ , and shift  $h$ . Each of those is limited to  $[0, 100]$  as follows:

$$a_t = (\Delta v_t, \Delta g_t, \Delta h_t), \quad (1)$$

$$v_{t+1} = \max(0, \min(100, v_t + d^v \Delta v_t)), \quad (2)$$

$$g_{t+1} = \max(0, \min(100, g_t + d^g \Delta g_t)), \quad (3)$$

$$h_{t+1} = \max(0, \min(100, h_t + d^h \Delta h_t)), \quad (4)$$

with scaling factors  $d^v = 1$ ,  $d^g = 10$ , and  $d^h = 5.75$ . The step size for changing shift is calculated as  $d^h = 20 \sin(15^\circ)/0.9 \approx 5.75$ .

After applying action  $a_t$ , the environment transitions to the next time step  $t + 1$  in which it enters internal state  $s_{t+1}$ . State  $s_t$  and successor state  $s_{t+1}$  are the Markovian states of the environment that are only partially observable to the agent.

An observable variable of the IB, setpoint  $p$ , influences the dynamical behavior of the environment but can never be changed by actions. An analogy to such a setpoint is, for example, the demanded load in a power plant or the wind speed actuating a wind turbine. As we will see in the upcoming description, different values of setpoint  $p$  will induce significant changes to the dynamics and stochasticity of the benchmark. The IB has two modes of operation: a) fixing the setpoint to a value  $p = \text{const}$ , thereby acting as a hyperparameter or b) as a time-varying external driver, making the dynamics become highly non-stationary. We give a detailed description of setting b) in Subsection III-D.

The set of observable state variables is completed by two reward relevant variables, consumption  $c_t$  and fatigue  $f_t$ . In the general RL setting a reward  $r_{t+1}$  for each transition  $t \rightarrow t + 1$  from state  $s_t$  via action  $a_t$  to the successor state  $s_{t+1}$  is drawn from a probability distribution depending on  $s_t$ ,  $a_t$ , and  $s_{t+1}$ . In the IB, the reward is given by a deterministic function of the successor state  $r_{t+1} = r(s_{t+1})$ , i.e.,

$$r_{t+1} = -c_{t+1} - 3f_{t+1}. \quad (5)$$

In the real-world tasks that motivated the IB, the reward function has always been known explicitly. In some cases it itself was subject to optimization and had to be adjusted to properly express the optimization goal. For the IB we therefore assume that the reward function is known and all variables influencing it are observable.

Thus the observation vector  $o_t$  at time  $t$  comprises current values of the set of observable state variables, which is a subset of all the variables of Markovian state  $s_t$ , i.e.,

- 1) the current steerings, velocity  $v_t$ , gain  $g_t$ , and shift  $h_t$ ,
- 2) the external driver, setpoint  $p_t$ ,
- 3) and the reward relevant variables consumption  $c_t$  and fatigue  $f_t$ .

Appendix A gives a complete overview on the IB's state space.

The data base for learning comprises tuples  $(o_t, a_t, o_{t+1}, r_t)$ . The agent is allowed to use all previous observation vectors and actions to estimate the Markovian state  $s_t$ .

The dynamics can be decomposed into three different sub-dynamics named operational cost, mis-calibration, and fatigue.

#### A. Dynamics of operational cost

The sub-dynamics of operational cost are influenced by the external driver setpoint  $p$  and two of the three steerings, velocity  $v$  and gain  $g$ . The current operational cost  $\theta_t$  is calculated as

$$\theta_t = \exp\left(\frac{2p_t + 4v_t + 2.5g_t}{100}\right). \quad (6)$$

The observation of  $\theta_t$  is delayed and blurred by the following convolution:

$$\theta_t^c = \frac{1}{9}\theta_{t-5} + \frac{2}{9}\theta_{t-6} + \frac{3}{9}\theta_{t-7} + \frac{2}{9}\theta_{t-8} + \frac{1}{9}\theta_{t-9}. \quad (7)$$

The convoluted operational cost  $\theta_t^c$  cannot be observed directly, instead it is modified by the second sub-dynamic, called mis-calibration, and finally subject to observation noise. The motivation for this dynamical behavior is that it is non-linear, it depends on more than one influence, and it is delayed and blurred. All those effects have been observed in industrial applications, like the heating process observable during combustion. In Figure 4b we give an example trajectory of the convolution process over a rollout of 200 time steps. The delayed and blurred relations between operational cost  $\theta_t$  and the convoluted costs  $\theta_t^c$  are clearly visible.

#### B. Dynamics of mis-calibration

The sub-dynamics of mis-calibration are influenced by external driver setpoint  $p$  and steering shift  $h$ . The goal is to reward an agent to oscillate in  $h$  in a pre-defined frequency around a specific operation point determined by setpoint  $p$ . Thereby, the reward topology is inspired by an example from quantum physics, namely Goldstone's "Mexican hat" potential.

In the first step, setpoint  $p$  and shift  $h$  are combined to an effective shift  $h^e$  calculated by:

$$h^e = \max\left(-1.5, \min\left(1.5, \frac{h}{20} - \frac{p}{50} - 1.5\right)\right). \quad (8)$$

Effective shift influences three latent variables, which are domain  $\delta$ , response  $\psi$ , and direction  $\phi$ . Domain  $\delta$  can enter two discrete states, which are *negative* and *positive*, represented by integer values  $-1$  and  $+1$ , respectively. Response  $\psi$  can enter two discrete states, which are *disadvantageous* and *advantageous*, represented by integer values  $-1$ , and  $+1$ , respectively. Direction  $\phi \in \{-6, -5, \dots, 6\}$  is a discrete index variable, yielding the position of the current optima in the mis-calibration penalty space.

Figure 2 is a visualization of the mis-calibration dynamics introduced in equation form in the following paragraphs. In each time step  $t \rightarrow t + 1$  the mis-calibration dynamics are transitioned starting with  $\delta$  and  $\psi$  as follows:

$$\hat{\delta}_{t+1} = \begin{cases} \delta_t & \text{if } |h^e| \leq z \\ \text{sgn}(h^e) & \text{else,} \end{cases} \quad (9)$$

$$\hat{\psi}_{t+1} = \begin{cases} 1 & \text{if } \delta_t \neq \hat{\delta}_{t+1} \\ \psi_t & \text{else,} \end{cases} \quad (10)$$

where safe zone  $z$  (area in the center of Figure 2) is calculated using  $z = \sin(\pi \cdot 15/180)/2 \approx 0.1309$ . Note that the policy itself is allowed to decide when to leave the safe zone.

In the next step, direction index  $\phi$  is updated accordingly:

$$\hat{\phi}_{t+1} = \phi_t + \Delta\phi_{t+1}, \text{ with} \quad (11)$$




Fig. 2. Visual description of the mis-calibration dynamics. Blue color represents areas of low penalty (-1.00), while yellow color represents areas of high penalty (1.23). If the policy keeps  $h^e$  in the so-called safe zone,  $\phi$  is driven towards 0 stepwise. When  $\phi = 0$  is reached the mis-calibration dynamics are reset, i.e., domain  $\delta = 1$  and response  $\psi = 1$ . The policy is allowed to start the rotation cycle at any time by leaving the safe zone and entering the *positive* or the *negative* domain. Consider *positive* domain  $\delta = 1$ : After initially leaving the safe zone, response is in the state *advantageous*, i.e.,  $\phi$  is increased stepwise. The upper right area is a reversal point for  $\phi$ . As soon as  $\phi = 6$ , response switches from *advantageous*  $\psi = 1$  to *disadvantageous*  $\psi = -1$ . In the subsequent time steps  $\phi$  is decreased until either the policy brings  $h^e$  back to the safe zone or  $\phi$  reaches the left boundary at -6. If the latter occurs,  $\phi$  is kept constant at -6, i.e., the policy yields a high penalty in each time step. Since the mis-calibration dynamics are symmetric around  $(\phi, h^e) = (0, 0)$ , opposite dynamics are applied in the *negative* domain at the lower part of the plot.

$$\Delta\phi_{t+1} = \begin{cases} -\text{sgn}(\phi_t) & \text{if } |h^e| \leq z \\ 0 & \text{if } |h^e| > z \wedge \phi_t = -6\delta_{t+1} \\ \hat{\psi}_{t+1} \cdot \text{sgn}(h^e) & \text{else.} \end{cases} \quad (12)$$

The first option realizes the return of  $\phi$  if the policy returns into the safe zone. The second option stops the rotation if  $\phi$  reaches the opposite domain bound (upper left and lower right area in Figure 2). The third option implements the cyclic movement of  $\phi$  depending on response  $\psi$  and the direction of effective shift  $h^e$ .

If, after this update, the absolute value of direction index  $\phi$  reaches or exceeds the predefined maximum index of 6 (upper right and lower left area in Figure 2), response enters state *disadvantageous* and index  $\phi$  is turned towards 0.

$$\hat{\psi}_{t+1} = \begin{cases} -1 & \text{if } |\hat{\phi}_{t+1}| \geq 6 \\ \hat{\psi}_{t+1} & \text{else.} \end{cases} \quad (13)$$

$$\phi_{t+1} = \begin{cases} 12 - ((\hat{\phi}_{t+1} + 24) \bmod 24) & \text{if } |\hat{\phi}_{t+1}| \geq 6 \\ \hat{\phi}_{t+1} & \text{else.} \end{cases} \quad (14)$$

In the final step of the mis-calibration state transition, it is checked if effective shift  $h^e$  has returned to safe zone  $z$  while at the same time direction index  $\phi$  has completed a full cycle (reset area in the center of Figure 2). If this is the case, domain  $\delta$  and response  $\psi$  are reset to their initial states *positive* and *advantageous*, respectively:

$$\delta_{t+1} = \begin{cases} 1 & \text{if } \phi_{t+1} = 0 \wedge |h^e| \leq z \\ \hat{\delta}_{t+1} & \text{else.} \end{cases} \quad (15)$$

$$\psi_{t+1} = \begin{cases} 1 & \text{if } \phi_{t+1} = 0 \wedge |h^e| \leq z \\ \hat{\psi}_{t+1} & \text{else.} \end{cases} \quad (16)$$

Note that in this state the policy can again decide to start a new cycle (positive or negative direction) or to remain in state  $\phi = 0$ .

The penalty landscape of mis-calibration is computed as follows. Based on the current value of  $\phi$ , the penalty function  $m(\phi, h^e)$  computes the performance of maintaining shift in the beneficial area. The penalty function  $m$  is defined as a linearly biased Goldstone potential computed by

$$m = -\alpha\omega^2 + \beta\omega^4 + \kappa\rho^s\omega. \quad (17)$$

The definition of radius  $\omega$  can be found in Appendix B. From direction index  $\phi$  the sine of direction angle  $\rho$  is calculated as follows:

$$\rho^s = \sin\left(\frac{\pi}{12}\phi\right). \quad (18)$$

Note that this sine function represents the optimal policy for the mis-calibration dynamics. Exemplary policy trajectories through the penalty landscape of mis-calibration are depicted and described in Figure 3.

The resulting mis-calibration  $m_t$  is added to the convoluted operational cost  $\theta_t^c$ , giving  $\hat{c}_t$ ,

$$\hat{c}_t = \theta_t^c + 25m_t. \quad (19)$$

Before being observable as consumption  $c_t$ , the modified operational cost  $\hat{c}_t$  is subject to heteroscedastic observation noise

$$c_t = \hat{c}_t + \text{gauss}(0, 1 + 0.02\hat{c}_t), \quad (20)$$

i.e., a Gaussian noise with zero mean and a standard deviation of  $\sigma = 1 + 0.02\hat{c}_t$ . In Figure 4c we show in an example rollout of 200 steps how both convoluted operational cost  $\theta_t^c$  and mis-calibration  $m_t$  affect consumption  $c_t$ .

### C. Dynamics of fatigue

The sub-dynamics of fatigue are influenced by the same variables as the sub-dynamics of operational cost, i.e., setpoint  $p$ , velocity  $v$ , and gain  $g$ . The IB is designed in such a way that, when changing the steering velocity  $v$  and gain  $g$  as to reduce the operational cost, fatigue will be increased,

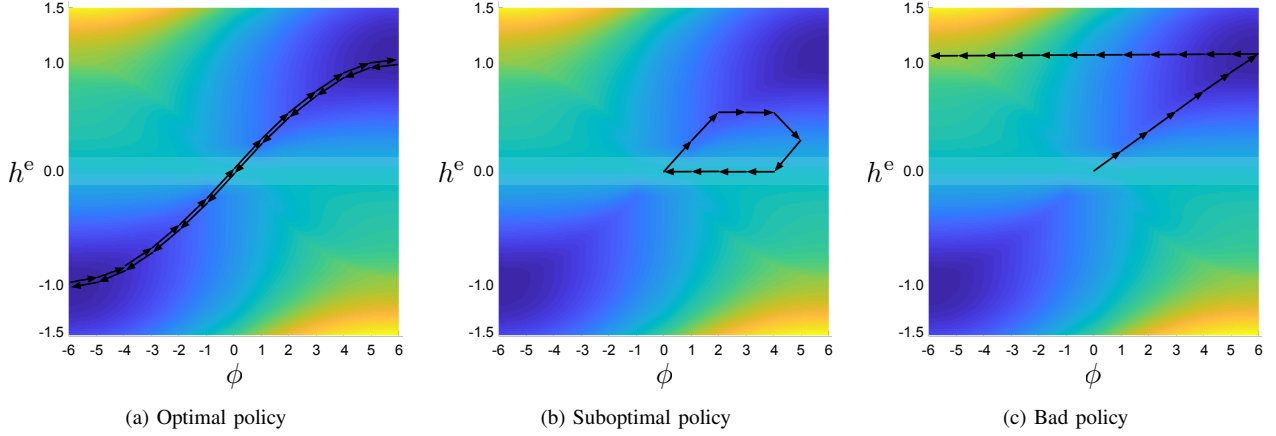


Fig. 3. Comparison of three mis-calibration policies. Depicted is a visual representation of the Goldstone potential based function  $m(\phi, h^e)$ . Areas yielding high penalty are colored yellow, while areas yielding low penalty are colored blue. The highlighted area in the center depicts the safe zone from  $-z$  to  $z$ . (a) A policy which maintains  $h^e$  such that a sine-shaped trajectory is generated yields lowest penalty. Note that the policy itself starts the rotation cycle at any time by leaving the safe zone. After returning to the safe zone, while at the same time  $\phi = 0$ , the dynamics are reset and a new cycle can be initiated at any following time step in positive or negative direction. (b) The depicted policy starts initiating the rotation cycle by leaving the safe zone, but returns after six steps. After this return  $\phi$  is decreased in four steps back to 0. Subsequently, the dynamics are reset. This policy yields lower penalty compared to a constant policy that remains in the safe zone the whole time. (c) The depicted policy approaches one of the global optima of  $m(\phi, h^e)$  by directly leaving the safe zone  $z$  by constantly increasing  $h^e$ . Subsequently, it remains at this point. However, the rotation dynamic yields a steady decrease in  $\phi$  after reaching the right boundary at  $\phi = 6$ . This decrease "pushes" the agent to the left, i.e., the penalties received are increased from step to step. After reaching the left boundary at  $\phi = -6$ , the dynamics remain in this area of high penalty. Note that the policy could bring the dynamics back to the initial state  $\phi = 0$  by returning to  $h^e < z$ . This benchmark property ensures that the best constant policies are the ones which remain in the safe zone.

leading to the desired multi-criterial task, with two reward-components showing opposite dependencies on the actions. The basic fatigue  $f^b$  is computed as

$$f^b = \max \left( 0, \frac{30000}{5v + 100} - 0.01g^2 \right). \quad (21)$$

From basic fatigue  $f^b$ , fatigue  $f$  is calculated by

$$f = f^b(1 + 2\alpha)/3, \quad (22)$$

where  $\alpha$  is an amplification. The amplification depends on two latent variables  $\mu^v$  and  $\mu^g$ , effective velocity  $v^e$ , and effective gain  $g^e$ . Furthermore, it is affected by noise,

$$\alpha = \begin{cases} \frac{1}{1 + \exp(-\text{gauss}(2.4, 0.4))} & \text{if } \max(\mu^v, \mu^g) = 1.2 \\ \max(\eta^v, \eta^g) & \text{else.} \end{cases} \quad (23)$$

In Eq. (23) we see that  $\alpha$  can undergo a bifurcation if one of the latent variables  $\mu^v$  or  $\mu^g$  reaches a value of 1.2. In that case,  $\alpha$  will increase and lead to higher fatigue, affecting the reward negatively.

The noise components  $\eta^v$  and  $\eta^g$ , as well as the latent variables  $h^v$  and  $h^g$ , depend on effective velocity  $v^e$ , and effective gain  $g^e$ . These are calculated by setpoint-dependent transformation functions

$$T^v(v, g, p) = \frac{g + p + 2}{v - p + 101}, \quad (24)$$

$$T^g(g, p) = \frac{1}{g + p + 1}. \quad (25)$$

Based on these transformation functions, effective velocity  $v^e$  and effective gain  $g^e$  are computed as follows:

$$v^e = \frac{T^v(v, g, p) - T^v(0, 100, p)}{T^v(100, 0, p) - T^v(0, 100, p)} \quad (26)$$

$$g^e = \frac{T^g(g, p) - T^g(100, p)}{T^g(0, p) - T^g(100, p)}. \quad (27)$$

To compute the noise components  $\eta^v$  and  $\eta^g$ , six random numbers are drawn from different random distributions:  $\eta^{ve}$  and  $\eta^{ge}$  are obtained by first sampling from an exponential distribution with mean 0.05 and applying the logistic function to these samples afterwards,  $\eta^{vb}$  and  $\eta^{gb}$  are drawn from binomial distributions  $\text{Binom}(1, v^e)$  and  $\text{Binom}(1, g^e)$ , respectively,  $\eta^{vu}$  and  $\eta^{gu}$  are drawn from a uniform distribution in  $[0, 1]$ . Noise components  $\eta^v$  and  $\eta^g$  are computed as follows:

$$\eta^v = \eta^{ve} + (1 - \eta^{ve})\eta^{vu}\eta^{vb}v^e \quad (28)$$

$$\eta^g = \eta^{ge} + (1 - \eta^{ge})\eta^{gu}\eta^{gb}g^e. \quad (29)$$

The latent variables  $\mu^v$  and  $\mu^g$  are calculated as

$$\mu_t^v = \begin{cases} v^e & \text{if } v^e \leq 0.05 \\ \min(5, 1.1\mu_{t-1}^v) & \text{if } v^e > 0.05 \wedge \mu_{t-1}^v \geq 1.2 \\ 0.9\mu_{t-1}^v + \frac{\eta^v}{3} & \text{else,} \end{cases} \quad (30)$$

$$\mu_t^g = \begin{cases} g^e & \text{if } g^e \leq 0.05 \\ \min(5, 1.1\mu_{t-1}^g) & \text{if } g^e > 0.05 \wedge \mu_{t-1}^g \geq 1.2 \\ 0.9\mu_{t-1}^g + \frac{\eta^g}{3} & \text{else.} \end{cases} \quad (31)$$

The sub-dynamic of fatigue results in a value for fatigue  $f$ , which is relevant for the reward function (Eq. 5). An

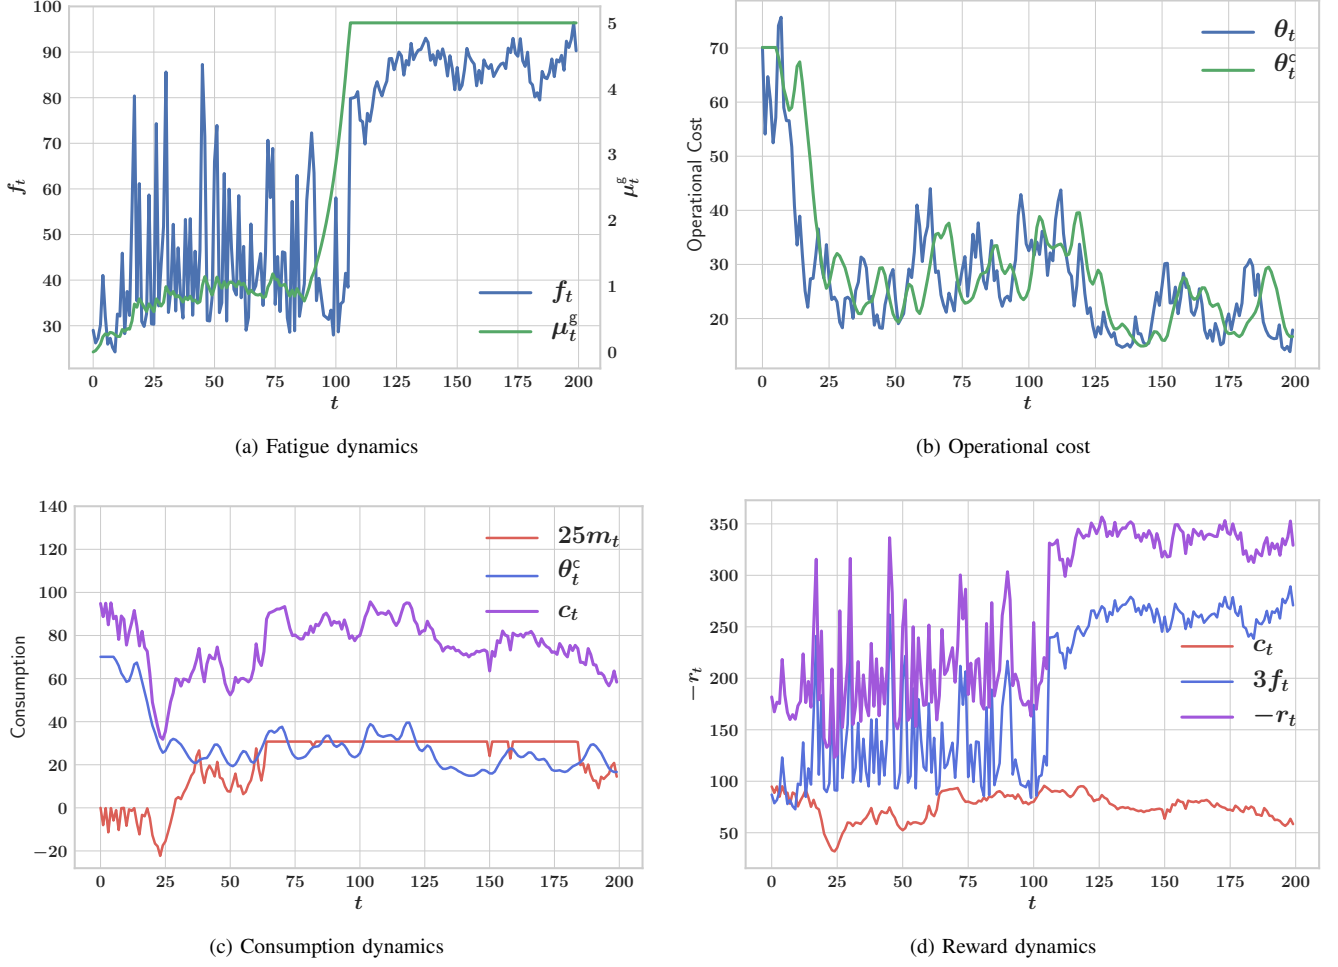


Fig. 4. Visualization of relevant variables of the IB in a rollout using random actions over 200 time steps. (a): Shown are latent variable  $\mu^g$  and fatigue  $f$ . As seen in Eq. (23), the latent variable can lead to a bifurcation of the dynamics. In the scenario shown at  $t = 90$ , we observe the beginning of a runaway effect that originates from the second case in Eq. (31). (b): Shown are operational cost  $\theta_t$  and convoluted sigma  $\theta_t^c$  given by Eq. (7). At around  $t = 10$  the delayed effect of the convolution is clearly visible:  $\theta_t$  decreases sharply while  $\theta_t^c$  is still ascending. (c): Shown is the composition of visible consumption  $c(t)$  (purple) by the two components  $\sigma_t^c$  and mis-calibration  $m_t$ . (d): Shown is the composition of final negative reward  $-r_t$  by its two components, fatigue (blue) and consumption (red). In this case, the runaway effect from Figure (a) has the most prominent effect on the reward signal.

example interplay of the components of the fatigue dynamics is visualized in Figure 4a. From  $t = 0$  up to  $t = 80$  we see the effect of the noise components described in Eq. (28): the combination of binomial and exponential noise components yields heterogeneous spike-like behavior. From  $t = 80$  to  $t = 100$  we observe a self-amplifying process in  $\mu^g$ . This self-amplification originates from the second case of Eq. (31). At around  $t = 100$ , the fatigue dynamics rapidly change towards higher, less noisy regions. This change originates from the bifurcation in  $\alpha$  in Eq. (23), which we pointed out earlier.

#### D. Setpoint dynamics

Setpoint  $p$  can either be kept constant or it is variable over time. In the variable setting, it will change by a constant value  $b$  over a fixed period of  $l$  time steps in the benchmark dynamics. Afterwards, a new sequence length and change rate is determined.

We sample sequence length  $l$  uniformly from  $\mathcal{U}\{1, 100\}$  and draw rate  $b$  from a mixture of a uniform distribution  $\mathcal{U}(0, 1)$  and a delta distribution  $\delta(x)$  with weighting probabilities 0.9 and 0.1. For each time step  $t + 1$  we update the setpoint according to:

$$p_{t+1} = \max(0, \min(100, p_t + b_{t+1})), \quad (32)$$

$$b_{t+1} = \begin{cases} -b_t & \text{if } (p_t = 0 \vee p_t = 100) \wedge z < 0.5 \\ b_t & \text{else,} \end{cases} \quad (33)$$

where  $z \sim \mathcal{U}(0, 1)$  will flip change rate  $b$  with a probability of 50% if the setpoint reaches one of the two bounds at  $p = 0$  or  $p = 100$ . Note that the equation above produces piecewise linear functions of constant change. We visualized four example trajectories in Figure 5.






Fig. 5. Four example trajectories of the setpoint of the IB in a variable setpoint setting.

#### IV. EXPERIMENTAL PROTOTYPES

The IB aims at being realistic in the sense, that it includes a variety of aspects that we found to be vital in industrial applications. In this section we want to outline prototypes of experimental settings that include key aspects present in industrial applications.

##### A. Batch Reinforcement Learning

In this setting, we are given an initial batch of data  $D$  from an already-running system and are asked to find a better (ideally near-optimal) policy.

The learners task is therefore to return a policy  $\pi(s_t)$  that can be deployed on the system at hand, solely based on the information provided by the batch [9]. These scenarios are common in real-world industry settings where exploration is usually restricted to avoid possible damage to the system.

Two scenarios using the IB for batch RL experiments are described subsequently.

*Random exploration:* In this setting, we generate a batch of state transitions using a random behavior policy, for instance by sampling action proposals from a uniform distribution. Example instances of these settings can be found in [18] and [10].

In the cited examples, the benchmark is initialized for ten different setpoints  $p \in \{10, 20, \dots, 100\}$  with the latent variables in their default values and the three steering variables at 50 each. Then, for each setpoint  $p$  the behavior policy is applied on the benchmark for 1,000 time steps, resulting in a total of 10,000 recorded state transitions. This process can be repeated to study the performance using different batch sizes.

For evaluation, the system is either initialized to its start settings [10], or at a random place in state space [18], at which point the policy drives the system autonomously.

*Safe behavior policy:* In real industrial settings, we seldom will run a fully random policy on the system at hand. A more realistic setting is that we have a batch of data generated by a safe, but suboptimal, behavior policy  $\pi_b$  with limited

randomness. In this setting, the task is to improve  $\pi_b$ . Unlike in the random exploration setting, the difficulty here is that large parts of the state space will be unavailable in the batch. The batch of data will likely contain more information of specific areas in state space and few information everywhere else. An example experiment can be found in [19].

##### B. Transfer Learning

A common situation in industrial control is that we have data from different industrial systems, or data from one industrial system that operates in different contexts. We expect that each instance will behave similarly on a global level, while we can expect significant deviations on a low level.

In the IB, this is realized by the setpoint  $p$ , a hyperparameter of the dynamics. Each value of  $p \in [0, \dots, 100]$  will define a different stochastic system, where the dissimilarity of two systems grows with the distance in  $p$ .

In transfer learning, we want to transfer our knowledge from system A to system B. For example, suppose we have a large batch  $D_1$  of state transitions from the IB with  $p = 50$ . We also have a small batch of state transition  $D_2$  with  $p = 75$ . If our goal is to learn a good model for a system with  $p = 75$ , the challenge of transfer learning is how to efficiently incorporate the batch  $D_1$  to improve learning. An example instance of this setup, albeit using pendulum dynamics, can be found in [20].

#### V. CONCLUSION

This paper introduced the IB, a novel benchmark for RL, inspired by industrial control. We have shown that it provides a useful addition to the set of existing RL benchmarks due to its unique combination of properties. Furthermore, we outlined prototype experimental setups relevant for industrial control. Our contributions are a step towards enabling other researchers to study RL in realistic industrial settings to expand the economical and societal impact of machine learning.

#### ACKNOWLEDGMENT

The project this report is based on was supported with funds from the German Federal Ministry of Education and Research under project number 01IB15001. The sole responsibility for the report's contents lies with the authors. The authors would like to thank Ludwig Winkler from TU Berlin for implementing the OpenAI Gym wrapper and sharing it with the community.

#### REFERENCES

- [1] M. Schlang, B. Feldkeller, B. Lang, P. T., and R. T. A., "Neural computation in steel industry," in *1999 European Control Conference (ECC)*, 1999, pp. 2922–2927.
- [2] T. A. Runkler, E. Gerstorfer, M. Schlang, E. Jünemann, and J. Hollatz, "Modelling and optimisation of a refining process for fibre board production," *Control engineering practice*, vol. 11, no. 11, pp. 1229–1241, 2003.
- [3] S. A. Hartmann and T. A. Runkler, "Online optimization of a color sorting assembly buffer using ant colony optimization," *Operations Research Proceedings 2007*, pp. 415–420, 2008.
- [4] A. M. Schaefer, D. Schneegass, V. Sterzing, and S. Udfluft, "A neural reinforcement learning approach to gas turbine control," in *2007 International Joint Conference on Neural Networks*, 2007, pp. 1691–1696.

- [5] A. Hans, D. Schneegass, A. M. Schaefer, and S. Udluft, "Safe exploration for reinforcement learning," in *2008 European Symposium on Artificial Neural Networks (ESANN)*, 2008, pp. 143–148.
- [6] M. G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling, "The arcade learning environment: An evaluation platform for general agents," *Journal of Artificial Intelligence Research*, vol. 47, pp. 253–279, 2013.
- [7] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2012, pp. 5026–5033.
- [8] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," *arXiv preprint arXiv:1509.02971*, 2015.
- [9] S. Lange, T. Gabel, and M. Riedmiller, "Batch reinforcement learning," in *Reinforcement Learning*. Springer, 2012, pp. 45–73.
- [10] D. Hein, S. Udluft, M. Tokic, A. Hentschel, T. A. Runkler, and V. Sterzing, "Batch reinforcement learning on the industrial benchmark: First experiences," in *2017 International Joint Conference on Neural Networks (IJCNN)*, 2017, pp. 4214–4221.
- [11] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press, 1998.
- [12] P. Abbeel, A. Coates, and A. Y. Ng, "Autonomous helicopter aerobatics through apprenticeship learning," *The International Journal of Robotics Research*, vol. 29, no. 13, pp. 1608–1639, 2010.
- [13] J. Randlev and P. Alström, "Learning to drive a bicycle using reinforcement learning and shaping," in *Proceedings of the Fifteenth International Conference on Machine Learning (ICML 1998)*, J. W. Shavlik, Ed. San Francisco, CA, USA: Morgan Kaufman, 1998, pp. 463–471.
- [14] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [15] H. van Seijen, M. Fatemi, J. Romoff, R. Larocche, T. Barnes, and J. Tsang, "Hybrid reward architecture for reinforcement learning," *arXiv preprint arXiv:1706.04208*, 2017.
- [16] 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. P. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis, "Mastering the game of Go with deep neural networks and tree search," *Nature*, vol. 529, no. 7587, pp. 484–489, 2016.
- [17] A. Hans and S. Udluft, "Efficient uncertainty propagation for reinforcement learning with limited data," in *2009 Proceedings of the International Conference on Artificial Neural Networks (ICANN 2009)*, 2009, pp. 70–79.
- [18] S. Depeweg, J. M. Hernández-Lobato, F. Doshi-Velez, and S. Udluft, "Learning and policy search in stochastic dynamical systems with Bayesian neural networks," *arXiv preprint arXiv:1605.07127*, 2016.
- [19] —, "Uncertainty decomposition in Bayesian neural networks with latent variables," *arXiv preprint arXiv:1706.08495*, 2017.
- [20] S. Spieckermann, S. Düll, S. Udluft, A. Hentschel, and T. A. Runkler, "Exploiting similarity in system identification tasks with recurrent neural networks," *Neurocomputing*, vol. 169, pp. 343–349, 2015.

## APPENDIX A STATE DESCRIPTION

Only a part of the state variables is observable. This observation vector is also called observable state, but one has to keep in mind, that it does not fulfill the Markov property. The observation vector  $o_t$  at time  $t$  comprises current values of velocity  $v_t$ , gain  $g_t$ , shift  $h_t$ , setpoint  $p_t$ , consumption  $c_t$ , and fatigue  $f_t$ .

The preferred minimal Markovian state fulfills the Markov property with the minimum number of variables. It comprises 20 values. These are the observation vector (velocity  $v_t$ , gain  $g_t$ , shift  $h_t$ , setpoint  $p_t$ , consumption  $c_t$ , and fatigue  $f_t$ ) plus some latent variables of the sub-dynamics. The sub-dynamics of operational cost add a list of previous operational costs,

$\theta_{t-i}$  with  $i \in 1, \dots, 9$ . Note that the current operational cost  $\theta_t$  is not part of this state definition. It would be redundant, as it can be calculated by  $v_t$ , gain  $g_t$ , and setpoint  $p_t$ . The sub-dynamics of mis-calibration need three additional latent variables,  $\delta$ ,  $\psi$ , and  $\phi$ , (Section III-B). The sub-dynamics of fatigue add 2 additional latent variables  $h^v$  and  $h^g$ , (Eq. (30) and (31)).

	text name or description	symbol
Observables	setpoint	$p_t$
	velocity	$v_t$
	gain	$g_t$
	shift	$h_t$
	consumption	$c_t$
	fatigue	$f_t$
Markovian state	operational cost at $t - 1$	$\theta_{t-1}$
	operational cost at $t - 2$	$\theta_{t-2}$
	operational cost at $t - 3$	$\theta_{t-3}$
	operational cost at $t - 4$	$\theta_{t-4}$
	operational cost at $t - 5$	$\theta_{t-5}$
	operational cost at $t - 6$	$\theta_{t-6}$
	operational cost at $t - 7$	$\theta_{t-7}$
	operational cost at $t - 8$	$\theta_{t-8}$
	operational cost at $t - 9$	$\theta_{t-9}$
	1 <sup>st</sup> latent variable of mis-calibration	$\delta$
	2 <sup>nd</sup> latent variable of mis-calibration	$\psi$
	3 <sup>rd</sup> latent variable of mis-calibration	$\phi$
	1 <sup>st</sup> latent variable fatigue	$\mu^v$
	2 <sup>nd</sup> latent variable fatigue	$\mu^g$

TABLE I  
IB MARKOVIAN STATE.

## APPENDIX B GOLDSTONE POTENTIAL BASED EQUATIONS

The resulting penalty of the mis-calibration reward component is computed by adopting a so-called linearly biased Goldstone potential. The following constants are pre-defined to subsequently compute the respective penalty:

$$\epsilon = \frac{\sqrt[3]{1 + \sqrt{2}}}{\sqrt{3}} \approx 0.7745, \quad (34)$$

$$\zeta = \epsilon + \frac{1}{3\epsilon} \approx 1.2048, \quad (35)$$

$$\lambda = 2\zeta^2 - \zeta^4 + 8\sqrt{\frac{2}{27}}\zeta \approx 3.4193, \quad (36)$$

$$\alpha = \frac{2}{\lambda} \approx 0.5849, \quad (37)$$

$$\beta = \frac{1}{\lambda} \approx 0.2924, \quad (38)$$

$$\kappa = \frac{-8\sqrt{\frac{2}{27}}}{\lambda} \approx -0.6367. \quad (39)$$

Given effective shift  $h^e$  and the sine of direction angle  $\rho$ , which is denoted as  $\rho^s$ , function  $\omega(\rho^s, h^e)$  is computed using



the following set of equations:

$$\omega = \begin{cases} (\frac{h^e |r^{\min}|}{|r^{\text{opt}}|}) & \text{if } |h^e| \leq |r^{\text{opt}}| \\ \text{sgn}(h^e) \cdot \hat{\omega} & \text{else,} \end{cases} \quad (40)$$

$$\hat{\omega} = |r^{\min}| + \frac{2 - |r^{\min}|}{(2 - |r^{\text{opt}}|)^{\hat{\omega}}} \cdot (|h^e| - |r^{\text{opt}}|)^{\hat{\omega}}, \quad (41)$$

$$\hat{\omega} = \frac{2 - |r^{\text{opt}}|}{2 - |r^{\min}|}, \quad (42)$$

$$r^{\text{opt}} = \varrho \cdot \max(|\rho^s|, 2z), \quad (43)$$

$$r^{\min} = \begin{cases} u + \frac{1}{3u} & \text{if } q < -\sqrt{1/27} \\ \varrho \sqrt{\frac{4}{3}} \cos\left(\frac{1}{3} \cdot \arccos(-q\sqrt{27})\right) & \text{else,} \end{cases} \quad (44)$$

$$u = \sqrt[3]{-\varrho q + \sqrt{q^2 - \frac{1}{27}}}, \quad (45)$$

$$q = \frac{\kappa |\rho^s|}{8\beta}, \quad (46)$$

$$\varrho = \text{sgn}(\rho^s). \quad (47)$$