
Munchausen Reinforcement Learning

Nino Vieillard
Google Research, Brain Team
Université de Lorraine, CNRS, Inria, IECL
F-54000 Nancy, France
vieillard@google.com

Olivier Pietquin
Google Research, Brain Team
pietquin@google.com

Matthieu Geist
Google Research, Brain Team
mfgeist@google.com

Abstract

Bootstrapping is a core mechanism in Reinforcement Learning (RL). Most algorithms, based on temporal differences, replace the true value of a transiting state by their current estimate of this value. Yet, another estimate could be leveraged to bootstrap RL: the current policy. Our core contribution stands in a very simple idea: adding the scaled log-policy to the immediate reward. We show that slightly modifying Deep Q -Network (DQN) in that way provides an agent that is competitive with distributional methods on Atari games, without making use of distributional RL, n -step returns or prioritized replay. To demonstrate the versatility of this idea, we also use it together with an Implicit Quantile Network (IQN). The resulting agent outperforms Rainbow on Atari, installing a new State of the Art with very little modifications to the original algorithm. To add to this empirical study, we provide strong theoretical insights on what happens under the hood – implicit Kullback-Leibler regularization and increase of the action-gap.

1 Introduction

Most Reinforcement Learning (RL) algorithms make use of Temporal Difference (TD) learning [29] in some ways. It is a well-known bootstrapping mechanism that consists in replacing the unknown true value of a transiting state by its current estimate and use it as a target for learning. Yet, agents compute another estimate while learning that could be leveraged to bootstrap RL: their current policy. Indeed, it reflects the agent’s hunch about which actions should be executed next and thus, which actions are good. Building upon this observation, our core contribution stands in a very simple idea: optimizing for the *immediate* reward *augmented* by the scaled log-policy of the agent when using any TD scheme. We insist right away that this is different from maximum entropy RL [36], that *subtracts* the scaled log-policy to *all* rewards, and aims at maximizing both the expected return and the expected entropy of the resulting policy. We call this general approach “*Munchausen Reinforcement Learning*” (M-RL), as a reference to a famous passage of *The Surprising Adventures of Baron Munchausen* by Raspe [24], where the Baron pulls himself out of a swamp by pulling on his own hair.

To demonstrate the genericity and the strength of this idea, we introduce it into the most popular RL agent: the seminal Deep Q -Network (DQN) [23]. Yet, DQN does not compute stochastic policies, which prevents using log-policies. So, we first introduce a straightforward generalization of DQN to maximum entropy RL [36, 17], and then modify the resulting TD update by adding the scaled log-policy to the immediate reward. The resulting algorithm, referred to as Munchausen-DQN (M-DQN), is thus genuinely a slight modification of DQN. Yet, it comes with strong empirical performances. On the Arcade Learning Environment (ALE) [6], not only it surpasses the original DQN by a large

margin, but it also overtakes C51 [8], the first agent based on distributional RL (distRL). As far as we know, M-DQN is the first agent not using distRL that outperforms a distRL agent¹. The current state of the art for single agent algorithms is considered to be Rainbow [18], that combines C51 with other enhancements to DQN, and does not rely on massively distributed computation (unlike R2D2 [19], SEED [12] or Agent57 [4]). To demonstrate the versatility of the M-RL idea, we apply the same recipe to modify Implicit Quantile Network (IQN) [11], a recent distRL agent. The resulting Munchausen-IQN (M-IQN) surpasses Rainbow, installing a new state of the art.

To support these empirical results, we provide strong theoretical insights about what happens under the hood. We rewrite M-DQN under an abstract dynamic programming scheme and show that it implicitly performs Kullback-Leibler (KL) regularization between consecutive policies. M-RL is not the first approach to take advantage of KL regularization [27, 2], but we show that, because this regularization is implicit, it comes with stronger theoretical guarantees. From this, we link M-RL to Conservative Value Iteration (CVI) [20] and Dynamic Policy Programming (DPP) [3] that were not introduced with deep RL implementations. We also draw connections with Advantage Learning (AL) [5, 7] and study the effect of M-RL on the action-gap [13]. While M-RL is not the first scheme to induce an increase of the action-gap [7], it is the first one that allows quantifying this increase.

2 Munchausen Reinforcement Learning

RL is usually formalized within the Markov Decision Processes (MDP) framework. An MDP models the environment and is a tuple $\{\mathcal{S}, \mathcal{A}, P, r, \gamma\}$, with \mathcal{S} and \mathcal{A} the state and action spaces, P the Markovian transition kernel, r the bounded reward function and γ the discount factor. The RL agent interacts with the MDP using a policy π , that associates to every state either an action (deterministic policy) or a distribution over actions (stochastic policy). The quality of this interaction is quantified by the expected discounted cumulative return, formalized as the state-action value function, $q_\pi(s, a) = \mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) | s_0 = s, a_0 = a]$, the expectation being over trajectories induced by the policy π and the dynamics P . An optimal policy satisfies $\pi_* \in \operatorname{argmax}_\pi q_\pi$. The associated optimal value function $q_* = q_{\pi_*}$ satisfies the Bellman equation $q_*(s, a) = r(s, a) + \gamma E_{s'|s, a}[\max_{a'} q_*(s', a')]$. A deterministic *greedy* policy satisfies $\pi(a|s) = 1$ for $a \in \operatorname{argmax}_{a'} q(s, a')$ and will be written $\pi \in \mathcal{G}(q)$. We also use *softmax* policies, $\pi = \operatorname{sm}(q) \Leftrightarrow \pi(a|s) = \frac{\exp q(s, a)}{\sum_{a'} \exp q(s, a')}$.

A standard RL agent maintains both a q -function and a policy (that can be implicit, for example $\pi \in \mathcal{G}(q)$), and it aims at learning an optimal policy. To do so, it often relies on Temporal Difference (TD) updates. To recall the principle of TD learning, we quickly revisit the classical Q -learning algorithm [34]. When interacting with the environment the agent observes transitions (s_t, a_t, r_t, s_{t+1}) . Would the optimal q -function q_* be known in the state s_{t+1} , the agent could use it as a learning target and build successive estimates as $q(s_t, a_t) \leftarrow q(s_t, a_t) + \eta(r_t + \gamma \max_{a'} q_*(s_{t+1}, a') - q(s_t, a_t))$, using the Bellman equation, η being a learning rate. Yet, q_* is unknown, and the agent actually uses its current estimate q instead, which is known as bootstrapping.

We argue that the q -function is not the sole quantity that could be used to bootstrap RL. Let's assume that an optimal deterministic policy π_* is known. The log-policy is therefore 0 for optimal actions, and $-\infty$ for sub-optimal ones. This is a very strong learning signal, that we could add to the reward to ease learning, without changing the optimal control. The optimal policy π_* being obviously unknown, we replace it by the agent's current estimate π , and we assume stochastic policies for numerical stability. To sum up, M-RL is a very simple idea, that consists in replacing r_t by $r_t + \alpha \ln \pi(a_t | s_t)$ in any TD scheme, assuming that the current agent's policy π is stochastic, so as to bootstrap the current agent's guess about what actions are good.

To demonstrate the generality of this approach, we use it to enhance the seminal DQN [23] deep RL algorithm. In DQN, the q -values are estimated by an online Q -network q_θ , with weights copied regularly to a target network $q_{\bar{\theta}}$. The agent behaves following a policy $\pi_\theta \in \mathcal{G}(q_\theta)$ (with ε -greedy exploration), and stores transitions (s_t, a_t, r_t, s_{t+1}) in a FIFO replay buffer \mathcal{B} . DQN performs

¹It appears that the benefits of distRL do not really come from RL principles, but rather from the regularizing effect of modelling a distribution and its role as an auxiliary task in a deep learning context [21].

stochastic gradient descent on the loss $\hat{\mathbb{E}}_{\mathcal{B}}[(q_{\bar{\theta}}(s_t, a_t) - \hat{q}_{\text{dqn}}(r_t, s_{t+1}))^2]$, regressing the target \hat{q}_{dqn} :

$$\hat{q}_{\text{dqn}}(r_t, s_{t+1}) = r_t + \gamma \sum_{a' \in \mathcal{A}} \pi_{\bar{\theta}}(a' | s_{t+1}) q_{\bar{\theta}}(s_{t+1}, a') \text{ with } \pi_{\bar{\theta}} \in \mathcal{G}(q_{\bar{\theta}}).$$

To derive Munchausen-DQN (M-DQN), we simply modify the regression target. M-RL assumes stochastic policies while DQN computes deterministic policies. A simple way to address this is to not only maximize the return, but also the entropy of the resulting policy, that is adopting the viewpoint of maximum entropy RL [36, 17]. It is straightforward to extend DQN to this setting, see Appx. A.1 for a detailed derivation. We call the resulting agent Soft-DQN (S-DQN). Let τ be the temperature parameter scaling the entropy, it just amounts to replace the original regression target by

$$\hat{q}_{\text{s-dqn}}(r_t, s_{t+1}) = r_t + \gamma \sum_{a' \in \mathcal{A}} \pi_{\bar{\theta}}(a' | s_{t+1}) \left(q_{\bar{\theta}}(s_{t+1}, a') - \tau \ln \pi_{\bar{\theta}}(a' | s_{t+1}) \right) \text{ with } \pi_{\bar{\theta}} = \text{sm}\left(\frac{q_{\bar{\theta}}}{\tau}\right), \quad (1)$$

where we highlighted the differences with DQN in blue. Notice that this is nothing more than the most straightforward discrete-actions version of Soft Actor-Critic (SAC) [17]. Notice also that in the limit $\tau \rightarrow 0$ we retrieve DQN. The last step to obtain M-DQN is to add the scaled log-policy to the reward. Let $\alpha \in [0, 1]$ be a scaling factor, the regression target of M-DQN is thus

$$\hat{q}_{\text{m-dqn}}(r_t, s_{t+1}) = r_t + \alpha \tau \ln \pi_{\bar{\theta}}(a_t | s_t) + \gamma \sum_{a' \in \mathcal{A}} \pi_{\bar{\theta}}(a' | s_{t+1}) \left(q_{\bar{\theta}}(s_{t+1}, a') - \tau \ln \pi_{\bar{\theta}}(a' | s_{t+1}) \right), \quad (2)$$

still with $\pi_{\bar{\theta}} = \text{sm}\left(\frac{q_{\bar{\theta}}}{\tau}\right)$, where we highlighted the difference with Soft-DQN in red (retrieved by setting $\alpha = 0$). Hence, M-DQN is genuinely obtained by replacing \hat{q}_{dqn} by $\hat{q}_{\text{m-dqn}}$ as the regression target of DQN. All details of the resulting algorithm are provide in Appx. B.1.




Figure 1: **Left:** Human-normalized mean scores. **Right:** Human-normalized median scores.

Despite being an extremely simple modification of DQN, M-DQN is very efficient. We show in Fig. 10 the Human-normalized mean and median scores for various agents on the full set of 60 Atari games of ALE (more details in Sec. 4). We observe that M-DQN significantly outperforms DQN, but also C51 [8]. As far we know, M-DQN is the first method that is not based on distRL which overtakes C51. These are quite encouraging empirical results.

To demonstrate the versatility of the M-RL principle, we also combine it with IQN [11], a recent and efficient distRL agent (note that IQN has had recent successors, such as Fully Parameterized Quantile Function (FQF) [35], to which in principle, we could also apply M-RL). We denote the resulting algorithm M-IQN. In a nutshell, IQN does not estimate the q -function, but the distribution of which the q -function is the mean, using a distributional Bellman operator. The (implicit) policy is still greedy according to the q -function, computed as the (empirical) mean of the estimated distribution. We apply the exact same recipe: derive soft-IQN using the principle of maximum entropy RL (which is as easy as for DQN), and add the scaled log-policy to the reward. For the sake of showing the generality of our method, we combine M-RL with a version of IQN that uses 3-steps returns (and we compare to IQN and Rainbow, that both use the same). We can observe on Fig. 10 that M-IQN outperforms Rainbow, both in terms of mean and median scores, and thus defines the new state of the art. In addition, even when using only 1-step returns, M-IQN still outperforms Rainbow. This result and the details of M-IQN can be found respectively in Appx. B.3 and B.1.

3 What happens under the hood?

The impressive empirical results of M-RL (see Sec. 4 for more) call for some theoretical insights. To provide them, we frame M-DQN in an abstract Approximate Dynamic Programming (ADP) framework and analyze it. We mainly provide two strong results: (1) M-DQN implicitly performs KL regularization between successive policies, which translates in an averaging effect of approximation errors (instead of accumulation in general ADP frameworks); (2) it increases the action-gap by a quantifiable amount which also helps dealing with approximation errors. We also use this section to draw connections with the existing literature in ADP. Let's first introduce some additional notations.

We write Δ_X the simplex over the finite set X and Y^X the set of applications from X to the set Y . With this, an MDP is $\{\mathcal{S}, \mathcal{A}, P \in \Delta_{\mathcal{S} \times \mathcal{A}}^{\mathcal{S} \times \mathcal{A}}, r \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}, \gamma \in (0, 1)\}$, the state and action spaces being assumed finite. For $f, g \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$, we define a component-wise dot product $\langle f, g \rangle = (\sum_a f(s, a)g(s, a))_s \in \mathbb{R}^{\mathcal{S}}$. This will be used with q -functions and (log-) policies, e.g. for expectations: $\mathbb{E}_{a \sim \pi(\cdot|s)}[q(s, a)] = \langle \pi, q \rangle(s)$. For $v \in \mathbb{R}^{\mathcal{S}}$, we have $Pv = (\mathbb{E}_{s'|s, a}[v(s')])_{s, a} = (\sum_{s'} P(s'|s, a)v(s'))_{s, a} \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$. We also defined a policy-induced transition kernel P_π as $P_\pi q = P\langle \pi, q \rangle$. With these notations, the Bellman evaluation operator is $T_\pi q = r + \gamma P_\pi q$ and its unique fixed point is q_π . An optimal policy still satisfies $\pi_* \in \operatorname{argmax}_{\pi \in \Delta_{\mathcal{A}}^{\mathcal{S}}} q_\pi$. The set of greedy policies can be written as $\mathcal{G}(q) = \operatorname{argmax}_{\pi \in \Delta_{\mathcal{A}}^{\mathcal{S}}} \langle \pi, q \rangle$. We'll also make use of the entropy of a policy, $\mathcal{H}(\pi) = -\langle \pi, \ln \pi \rangle$, and of the KL between two policies, $\operatorname{KL}(\pi_1 || \pi_2) = \langle \pi_1, \ln \pi_1 - \ln \pi_2 \rangle$.

A softmax is the maximizer of the Legendre-Fenchel transform of the entropy [9, 32], $\operatorname{sm}(q) = \operatorname{argmax}_{\pi} \langle \pi, q \rangle + \mathcal{H}(\pi)$. Using this and the introduced notations, we can write M-DQN in the following abstract form (each iteration consists of a greedy step and an evaluation step):

$$\begin{cases} \pi_{k+1} = \operatorname{argmax}_{\pi \in \Delta_{\mathcal{A}}^{\mathcal{S}}} \langle \pi, q_k \rangle + \tau \mathcal{H}(\pi) \\ q_{k+1} = r + \alpha \tau \ln \pi_{k+1} + \gamma P \langle \pi_{k+1}, q_k - \tau \ln \pi_{k+1} \rangle + \epsilon_{k+1}. \end{cases} \quad \text{M-VI}(\alpha, \tau) \quad (3)$$

We call the resulting scheme Munchausen Value Iteration, or M-VI(α, τ). The term ϵ_{k+1} stands for the error between the actual and the ideal update (sampling instead of expectation, approximation of q_k by a neural network, fitting of the neural network). Removing the **red term**, we retrieve approximate VI (AVI) regularized by a scaled entropy, as introduced by Geist et al. [15], of which Soft-DQN is an instantiation (as well as SAC, with additional error in the greedy step). Removing also the **blue term**, we retrieve the classic AVI [26], of which DQN is an instantiation.

To get some insights, we rewrite the evaluation step, setting $\alpha = 1$ and with $q'_k \triangleq q_k - \tau \ln \pi_k$:

$$\begin{aligned} q_{k+1} &= r + \tau \ln \pi_{k+1} + \gamma P \langle \pi_{k+1}, q_k - \tau \ln \pi_{k+1} \rangle + \epsilon_{k+1} \\ \Leftrightarrow q_{k+1} - \tau \ln \pi_{k+1} &= r + \gamma P \langle \pi_{k+1}, q_k - \tau \ln \pi_k - \tau \ln \frac{\pi_{k+1}}{\pi_k} \rangle + \epsilon_{k+1} \\ \Leftrightarrow q'_{k+1} &= r + \gamma P (\langle \pi_{k+1}, q'_k \rangle - \tau \operatorname{KL}(\pi_{k+1} || \pi_k)) + \epsilon_{k+1}. \end{aligned}$$

Then, the greedy step can be rewritten as (looking at what π_{k+1} maximizes)

$$\langle \pi, q_k \rangle + \tau \mathcal{H}(\pi) = \langle \pi, q'_k + \tau \ln \pi_k \rangle - \tau \langle \pi, \ln \pi \rangle = \langle \pi, q'_k \rangle - \tau \operatorname{KL}(\pi || \pi_k). \quad (4)$$

We have just shown that M-VI(1, τ) implicitly performs KL regularization between successive policies.

This is a very insightful result as KL regularization is the core component of recent efficient RL agents such as TRPO [27] or MPO [2]. It is extensively discussed by Vieillard et al. [32]. Interestingly, we can show that the sequence of policies produced by M-VI(α, τ) is the same as the one of their Mirror Descent VI (MD-VI), with KL scaled by $\alpha\tau$ and entropy scaled by $(1 - \alpha)\tau$. Thus, M-VI(α, τ) is equivalent to MD-VI($\alpha\tau, (1 - \alpha)\tau$), as formalized below (proof in Appx. A.2).

Theorem 1. For any $k \geq 0$, define $q'_k = q_k - \alpha\tau \ln \pi_k$, we have

$$(3) \Leftrightarrow \begin{cases} \pi_{k+1} = \operatorname{argmax}_{\pi \in \Delta_{\mathcal{A}}^{\mathcal{S}}} \langle \pi, q'_k \rangle - \alpha\tau \operatorname{KL}(\pi || \pi_k) + (1 - \alpha)\tau \mathcal{H}(\pi) \\ q'_{k+1} = r + \gamma P (\langle \pi_{k+1}, q'_k \rangle - \alpha\tau \operatorname{KL}(\pi_{k+1} || \pi_k) + (1 - \alpha)\tau \mathcal{H}(\pi_{k+1})) + \epsilon_{k+1} \end{cases}$$

Moreover, [32, Thm. 1] applies to M-VI(1, τ) and [32, Thm. 2] applies to M-VI($\alpha < 1, \tau$).

In their work, Vieillard et al. [32] show that using regularization can reduce the dependency to the horizon $(1 - \gamma)^{-1}$ and that using a KL divergence allows for a compensation of the errors ϵ_k over

iterations, which is not true for classical ADP. We refer to them for a detailed discussion on this topic. However, we would like to highlight that they acknowledge that *their theoretical analysis does not apply to the deep RL setting*. The reason being that their analysis does not hold when the greedy step is approximated, and they deem as impossible to do the greedy step exactly when using neural network. Indeed, computing π_{k+1} by maximizing eq. (4) yields an analytical solution proportional to $\pi_k \exp(\frac{q_k}{\tau})$, and that thus depends on the previous policy π_k . Consequently, the solution to this equation cannot be computed exactly when using deep function approximation (unless one would be willing to remember every computed policy). On the contrary, *their analysis applies in our deep RL setting*. In M-VI, the KL regularization is implicit, so *we do not introduce errors in the greedy step*. To be precise, the greedy step of M-VI is only a softmax of the q -function, which can be computed exactly in a discrete actions setting, even when using deep networks. Their strong bounds for MD-VI therefore hold for M-VI, as formalized in Thm. 1, and in particular for M-DQN.

Indeed, let $q_{\bar{\theta}_k}$ be the k^{th} update of the target network, write $q_k = q_{\bar{\theta}_k}$, $\pi_{k+1} = \text{sm}(\frac{q_k}{\tau})$, and define $\epsilon_{k+1} = q_{k+1} - (r + \alpha \ln \pi_{k+1} - \gamma P \langle \pi_{k+1}, q_k - \tau \ln \pi_{k+1} \rangle)$, the difference between the actual update and the ideal one. As a direct corollary of Thm. 1 and [32, Thm. 1], we have that, for $\alpha = 1$,

$$\|q_* - q_{\pi_k}\|_{\infty} \leq \frac{2}{1-\gamma} \left\| \frac{1}{k} \sum_{j=1}^k \epsilon_j \right\|_{\infty} + \frac{4}{(1-\gamma)^2} \frac{r_{\max} + \tau \ln |\mathcal{A}|}{k},$$

with r_{\max} the maximum reward (in absolute value), and with q_{π_k} the true value function of the policy of M-DQN. This is a very strong bound. The error term is $\|\frac{1}{k} \sum_{j=1}^k \epsilon_j\|_{\infty}$, to be compared to the one of AVI [26], $(1-\gamma) \sum_{j=1}^k \gamma^{k-j} \|\epsilon_j\|_{\infty}$. Instead of having a discounted sum of the norms of the errors, we have the norm of the average of the errors. This is very interesting, as it allows for a compensation of errors between iterations instead of an accumulation (sum and norm do not commute). The error term is scaled by $(1-\gamma)^{-1}$ (the average horizon of the MDP), while the one of AVI would be scaled by $(1-\gamma)^{-2}$. This is also quite interesting, a γ close to 1 impacts less negatively the bound. We refer to [32, Sec. 4.1] for further discussions about the advantage of this kind of bounds. Similarly, we could derive a bound for the case $\alpha < 1$, and even more general and meaningful component-wise bounds. We defer the statement of these bounds and their proofs to Appx. A.3.

From Eq. (3), we can also relate the proposed approach to another part of the literature. Still from basic properties of the Legendre-Fenchel transform, we have that $\max_{\pi} \langle q, \pi \rangle + \tau \mathcal{H}(\pi) = \langle \pi_{k+1}, q_k \rangle + \tau \mathcal{H}(\pi_{k+1}) = \ln \langle 1, \exp q \rangle$. In other words, if the maximizer is the softmax, the maximum is the log-sum-exp. Using this, Eq. (3) can be rewritten as (see Appx. A.4 for a detailed derivation)

$$q_{k+1} = r + \gamma P(\tau \ln \langle 1, \exp \frac{q_k}{\tau} \rangle) + \alpha(q_k - \tau \ln \langle 1, \exp \frac{q_k}{\tau} \rangle) + \epsilon_{k+1}. \quad (5)$$

This is very close to Conservative Value Iteration² (CVI) [20], a purely theoretical algorithm, as far as we know. With $\alpha = 0$ (without Munchausen), we get Soft Q-learning [14, 16]. Notice that with this, CVI can be seen as soft Q-learning plus a scaled and smooth advantage (the term $\alpha(q_k - \tau \ln \langle 1, \exp \frac{q_k}{\tau} \rangle)$). With $\alpha = 1$, we retrieve a variation of Dynamic Policy Programming (DPP) [3, Appx. A]. DPP has been extended to a deep learning setting [30], but it is less efficient than DQN³ [32]. Taking the limit $\tau \rightarrow 0$, we retrieve Advantage Learning (AL) [5, 7] (see Appx. A.4):

$$q_{k+1} = r + \gamma P \langle \pi_{k+1}, q_k \rangle + \alpha(q_k - \langle \pi_{k+1}, q_k \rangle) + \epsilon_{k+1} \text{ with } \pi_{k+1} \in \mathcal{G}(q_k). \quad (6)$$

AL aims at increasing the action-gap [13] defined as the difference, for a given state, between the (optimal) value of the optimal action and that of the suboptimal ones. The intuitive reason to want a large action-gap is that it can mitigate the undesirable effects of approximation and estimation errors made on q on the induced greedy policies. Bellemare et al. [7] have introduced a family of Bellman-like operators that are gap-increasing. Not only we show that M-VI is gap-increasing but we also quantify the increase. To do so, we introduce some last notations. As we explained before, with $\alpha = 0$, M-VI(0, τ) reduces to AVI regularized by an entropy (that is, maximum entropy RL). Without error, it is known that the resulting regularized MDP has a unique optimal policy π_*^{τ} and a unique optimal q -function⁴ q_*^{τ} [15]. This being defined, we can state our result (proven in Appx. A.5).

²In CVI, $\langle 1, \exp \frac{q_k}{\tau} \rangle$ is replaced by $\langle \frac{1}{|\mathcal{A}|}, \exp \frac{q_k}{\tau} \rangle$.

³In fact, Tsurumine et al. [30] show better performance for deep DPP than for DQN in their setting. Yet, their experiment involves a small number of interactions, while the function estimated by DPP is naturally diverging. See [33, Sec. 6] for further discussion about this.

⁴It can be related to the unregularized optimal q -function, $\|q_*^{\tau} - q_*\|_{\infty} \leq \frac{\tau \ln |\mathcal{A}|}{1-\gamma}$ [15].

Theorem 2. For any state $s \in \mathcal{S}$, define the action-gap of an MPD regularized by an entropy scaled by τ as $\delta_*^\tau(s) = \max_a q_*^\tau(s, a) - q_*^\tau(s, \cdot) \in \mathbb{R}_+^A$. Define also $\delta_k^{\alpha, \tau}(s)$ as the action-gap for the k^{th} iteration of M-VI(α, τ), without error ($\epsilon_k = 0$): $\delta_k^{\alpha, \tau}(s) = \max_a q_k(s, a) - q_k(s, \cdot) \in \mathbb{R}_+^A$. Then, for any $s \in \mathcal{S}$, for any $0 \leq \alpha \leq 1$ and for any $\tau > 0$, we have

$$\lim_{k \rightarrow \infty} \delta_k^{\alpha, \tau}(s) = \frac{1 + \alpha}{1 - \alpha} \delta_*^{(1-\alpha)\tau}(s),$$

with the convention that $\infty \cdot 0 = 0$ for $\alpha = 1$.

Thus, the original action-gap is multiplied by $\frac{1+\alpha}{1-\alpha}$ with M-VI. In the limit $\alpha = 1$, it is even infinite (and zero for the optimal actions). This suggests choosing a large value of α , but not too close to 1 (for numerical stability: if having a large action-gap is desirable, having an infinite one is not).

4 Experiments

Munchausen agents. We implement M-DQN and M-IQN as variations of respectively DQN and IQN from Dopamine [10]. We use the same hyperparameters for IQN⁵, and we only change the optimizer from RMSProp to Adam for DQN. This is actually not anodyne, and we study its impact in an ablation study. We also consider a Munchausen-specific modification, *log-policy clipping*. Indeed, the log-policy term is not bounded, and can cause numerical issues if the policy becomes too close to deterministic. Thus, with a hyperparameter $l_0 < 0$, we replace $\tau \ln \pi(a|s)$ by $[\tau \ln \pi(a|s)]_{l_0}^0$, where $[\cdot]_x^y$ is the clipping function. For numerical stability, we use a specific log-sum-exp trick to compute the log-policy (see App. B.1). Hence, we add three parameters to the modified agent: α , τ and l_0 . After some tuning on a few Atari games, we found a working zone for these parameters to be $\alpha = 0.9$, $\tau = 0.03$ and $l_0 = -1$, used for all experiments, in M-DQN and M-IQN. All details about the rest of the parameters can be found in Appx. B.1. DQN and IQN use ϵ -greedy policies to interact with the environment. Although M-DQN and M-IQN produce naturally stochastic policies, we use the same ϵ -greedy policies. We discuss this further in Appx. B.2, where we also compare to stochastic policies.

Baselines. First, we consider both DQN and IQN, as these are the algorithms we modify. Second, we compare to C51 because, as far as we know, it has never been outperformed by a non-distRL agent before. We also consider Rainbow, as it stands for being the state-of-the-art non-distributed agent on ALE. All our baselines are taken from Dopamine. For Rainbow, this version doesn’t contain all the original improvements, but only the ones deemed as the more important and efficient by Hessel et al. [18]: n -steps returns and Prioritized Experience Replay (PER) [25], on top of C51.

Task. We evaluate our methods and the baselines in the ALE environment, *i.e.* on the full set of 60 Atari games. Notice that it is not a “canonical” environment. For example, choosing to end an episode when an agent loses a life or after game-over can dramatically change the score an agent can reach (*e.g.*, [10, Fig. 4]). The same holds for using sticky actions, introducing stochasticity in the dynamics (*e.g.*, [10, Fig. 6]). Even the ROMs could be different, with unpredictable consequences (*e.g.* different video encoding). Here, we follow the methodological best practices proposed by Machado et al. [22] and instantiated in Dopamine [10], that also makes the ALE more challenging. Notably, the results we present are hardly comparable to the ones presented in the original publications of DQN [23], C51 [8], Rainbow [18] or IQN [11], that use a different, easier, setting. Yet, for completeness, we report results on one game (Asterix) using an ALE setting as close as possible to the original papers, in Appx. B.4: the baseline results match the previously published ones, and M-RL still raises improvement. We also highlight that we stick to a single-agent version of the environment: we do not claim that our method can be compared to highly distributed agents, such as R2D2 [19] or Agent57 [4], that use several versions of the environment in parallel, and train on a much higher number of frames (around 10G frames vs 200M here). Yet, we are confident that our approach could easily apply to such agents.

Metrics. All algorithms are evaluated on the same training regime (details in Appx.B.1), during 200M frames, and results are averaged over 3 seeds. As a metric for any games, we compute the “baseline-normalized” score, for each iteration (here, 1M frames), normalized so that 0% corresponds to a random score, and 100% to the final performance of the baseline. At each iteration, the score is

⁵By default, Dopamine’s IQN uses 3-steps returns. We rather consider 1-step returns, as in [11].

the undiscounted sum of rewards, averaged over the last 100 learning episodes. The normalized score is then $\frac{a-r}{|b-r|}$, with a the score of the agent, b the score of the baseline, and r the score of a random policy. For a human baseline, the scores are those provided in Table 3 (Appx. B.6), for an agent baseline the score is the one after 200M frames. With this, we provide aggregated results, showing the mean and the median over games, as learning proceeds when the baseline is the human score (e.g., Fig. 1), or after 200M steps with human and Rainbow baselines in Tab. 3 (more results in Appx. B.6, as learning proceeds). We also compute a “baseline-improvement” score as $\frac{a-b}{|b-r|}$, and use it to report a per-game improvement after 200M frames (Fig. 4, M-Agent versus Agent, or Appx. B.6).

Action-gap. We start by illustrating the action-gap phenomenon suggested by Thm. 2. To do so, let q_θ be the q -function of a given agent after training for 200M steps. At any time-step t , write $\hat{a}_t \in \operatorname{argmax}_{a \in \mathcal{A}} q_\theta(s_t, a)$ the current greedy action, we compute the empirical action-gap as the difference of estimated values between the best and second best actions, $q_\theta(s_t, \hat{a}_t) - \max_{a \in \mathcal{A} \setminus \{\hat{a}_t\}} q_\theta(s_t, a)$. We do so for M-DQN, for AL (that was introduced specifically to increase the action-gap) and for DQN with Adam optimizer (Adam DQN), as both build on top of it (only changing the regression targets, see Appx. B.1 for details).




Figure 2: Action-gaps (Asterix).

We consider the game Asterix, for which the final average performance of the agents are (roughly) 15k for Adam DQN, 13k for AL and 20k for M-DQN. We report the results on Fig. 2: we run each agent for 10 trajectories, and average the resulting action-gaps (the length of the resulting trajectory is the one of the shorter trajectory, we also apply an exponential smoothing of 0.99). Both M-DQN and AL increase the action-gaps compared to Adam DQN. If AL increases it more, it seems also to be less stable, and less proportional to the original action-gap. Despite this increase, it performs worse than Adam DQN (13k vs 15k), while M-DQN increases it and performs better (20k vs 15k). An explanation to this phenomenon could be the one of Van Seijen et al. [31], who suggest that what is important is not the value of the action gap itself, but its uniformity over the state-action space: here, M-DQN seems to benefit from a more stable action-gap than AL. This figure is for an illustrative purpose, one game is not enough to draw conclusions. Yet, the following ablation shows that globally M-DQN performs better than AL. Also, it benefits from more theoretical justifications (not only quantified action-gap increase, but also implicit KL-regularization and resulting performance bounds).

Ablation study. We’ve build M-DQN from DQN by adding the Adam optimizer (Adam DQN), extending it to maximum entropy RL (Soft-DQN, Eq. (1)), and then adding the Munchausen term (M-DQN, Eq. (2)). A natural ablation is to remove the Munchausen term, and use only maximum entropy RL, by considering M-DQN with $\alpha = 0$ (instead of 0.9 for M-DQN), and the same τ (here, $3e - 2$), which would give Soft-DQN(τ). However, Thm. 1 states that M-DQN performs entropy regularization with an implicit coefficient of $(1 - \alpha)\tau$, so to compare M-DQN and Soft-DQN fairly, one should evaluate Soft-DQN with such a temperature, that is $3e - 3$ in this case. We denote this ablation as Soft-DQN($(1 - \alpha)\tau$). As sketched in Sec. 3, AL can also be seen as a limit case (on an abstract way, as $\tau \rightarrow 0$, see also Appx. B.1 for details on the algorithm). We provide an ablation study of all these variations, all using Adam (except DQN), in Fig. 3. All methods perform better than DQN. Adam DQN performs very well and is even competitive with C51. This is an interesting insight, as changing the optimizer compared to the published parameters dramatically improves the performance, and Adam DQN could be considered as a better baseline⁶. Surprisingly, if better than DQN, Soft-DQN does not perform better than Adam DQN. This suggests that maximum entropy RL alone might not be sufficient. We kept the temperature $\tau = 0.03$, and one could argue that it was not tuned for Soft DQN, but it is on par with the temperature of similar algorithms [28, 32]. We observe that AL performs better than Adam DQN. Again, we kept $\alpha = 0.9$, but this is consistent with the best performing parameter of Bellemare et al. [7, e.g., Fig. 7]. The proposed M-DQN outperforms all other methods, both in mean and median, and especially Soft-DQN by a significant margin (the sole difference being the Munchausen term).

Comparison to the baselines. We report aggregated results as Human-normalized mean and median scores on Figure 1, that compares the Munchausen agents to the baselines. M-DQN is largely

⁶To be on par with the literature, we keep using the published DQN as the baseline for other experiments.




Figure 3: Ablation study of M-DQN: Human-normalized mean (**left**) and median (**right**) scores.

Table 1: Mean/median Human/Rainbow-normalized scores at 200M frames, on the 60 games, averaged over 3 random seeds. In **bold** are the best of each column, and in **blue** over Rainbow. We also provide the number of improved games (compared to Human and Rainbow).

	Human-normalized			Rainbow-normalized		
	Mean	Median	#Improved	Mean	Median	#Improved
M-DQN	340%	124%	37	89%	92%	21
M-IQN	563%	165%	43	130%	109%	38
RAINBOW	414%	150%	43	100%	100%	-
IQN	441%	139%	41	105%	99%	27
C51	339%	111%	33	84%	70%	11
DQN	228%	71%	23	51%	51%	3

over DQN, and outperforms C51 both in mean and median. It is remarkable that M-DQN, justified by theoretically sound RL principles and without using common deep RL tricks like n -steps returns, PER or distRL, is competitive with distRL methods. It is even close to IQN (in median), considered as the best distRL-based agent. We observe that M-IQN, that combines IQN with Munchausen principle, is better than all other baselines, by a significant margin in mean. We also report the final Human-normalized and Rainbow-normalized scores of all the algorithms in Table 1. These results are on par with the Human-normalized scores of Fig. 1 (see Appx. B.6 for results over frames). M-DQN is still close to IQN in median, is better than DQN, and C51, while M-IQN is the best agent w.r.t. all metrics.

Per-game improvements. In Figure 4, we report the improvement for each game of the Munchausen agents over the algorithms they modify. The “Munchausened” versions show significant improvements, on a large majority of Atari games (53/60 for M-DQN vs DQN, 40/60 for M-IQN vs IQN). This result also explains the sometime large difference between the mean and median metrics, as some games benefit from a particularly large improvement. All learning curves are in Appx B.6.

5 Conclusion

In this work, we presented a simple extension to RL algorithms: Munchausen RL. This method augments the immediate rewards by the scaled logarithm of the policy computed by an RL agent. We applied this method to a simple variation of DQN, Soft-DQN, resulting in the M-DQN algorithm. M-DQN shows large performance improvements: it outperforms DQN on 53 of the 60 Atari games, while simply using a modification of the DQN loss. In addition, it outperforms the seminal distributional RL algorithm C51. We also extended the Munchausen idea to distributional RL, showing that it could be successfully combined with IQN to outperform the Rainbow baseline. Munchausen-DQN relies on theoretical foundations. To show that, we have studied an abstract Munchausen Value Iteration scheme and shown that it implicitly performs KL regularization. Notably, the strong theoretical results of [32] apply to M-DQN. By rewriting it in an equivalent ADP form, we have related our approach to the literature, notably to CVI, DPP and AL. We have shown that M-VI increases the




Figure 4: Per-game improvement of M-DQN vs DQN (top) and of M-IQN vs IQN (bottom).

action-gap, and we have quantified this increase, that can be infinite in the limit. In the end, this work highlights that a thoughtful revisiting of the core components of reinforcement learning can lead to new and efficient deep RL algorithms.

Broader impact

The core contribution of this work is to propose a new RL algorithm, that surpasses state of the art results on a challenging discrete actions environment. We believe it can impact positively the RL community, as it shades light on fundamental ideas, justified by deep theoretical foundations, that proves to be efficient in practice. Outside of the RL community, the impact of this paper is part of the global impact of RL. This work is mainly algorithmic and theoretical, with no specific applications in mind, but participates to the general development of efficient and practical RL methods.

Funding transparency statement

Nothing to disclose.

References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [2] Abbas Abdolmaleki, Jost Tobias Springenberg, Yuval Tassa, Remi Munos, Nicolas Heess, and Martin Riedmiller. Maximum a posteriori policy optimisation. In *International Conference on learning Representations (ICLR)*, 2018.
- [3] Mohammad Gheshlaghi Azar, Vicenç Gómez, and Hilbert J Kappen. Dynamic policy programming. *Journal of Machine Learning Research*, 13(Nov):3207–3245, 2012.
- [4] Adrià Puigdomènech Badia, Bilal Piot, Steven Kapturowski, Pablo Sprechmann, Alex Vitvitskyi, Daniel Guo, and Charles Blundell. Agent57: Outperforming the atari human benchmark. *arXiv preprint arXiv:2003.13350*, 2020.
- [5] Leemon C Baird III. *Reinforcement Learning Through Gradient Descent*. PhD thesis, US Air Force Academy, US, 1999.
- [6] Marc G Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.
- [7] Marc G Bellemare, Georg Ostrovski, Arthur Guez, Philip S Thomas, and Rémi Munos. Increasing the action gap: New operators for reinforcement learning. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2016.
- [8] Marc G Bellemare, Will Dabney, and Rémi Munos. A distributional perspective on reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2017.
- [9] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [10] Pablo Samuel Castro, Subhodeep Moitra, Carles Gelada, Saurabh Kumar, and Marc G Bellemare. Dopamine: A research framework for deep reinforcement learning. *arXiv preprint arXiv:1812.06110*, 2018.
- [11] Will Dabney, Georg Ostrovski, David Silver, and Rémi Munos. Implicit quantile networks for distributional reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2018.

- [12] Lasse Espeholt, Raphael Marinier, Piotr Stanczyk, Ke Wang, and Marcin Michalski. SEED RL: Scalable and Efficient Deep-RL with Accelerated Central Inference. In *International Conference on Learning Representations (ICLR)*, 2020.
- [13] Amir-massoud Farahmand. Action-gap phenomenon in reinforcement learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 172–180, 2011.
- [14] Roy Fox, Ari Pakman, and Naftali Tishby. Taming the noise in reinforcement learning via soft updates. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2016.
- [15] Matthieu Geist, Bruno Scherrer, and Olivier Pietquin. A Theory of Regularized Markov Decision Processes. In *International Conference on Machine Learning (ICML)*, 2019.
- [16] Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. Reinforcement Learning with Deep Energy-Based Policies. In *International Conference on Machine Learning (ICML)*, 2017.
- [17] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning (ICML)*, 2018.
- [18] Matteo Hessel, Joseph Modayil, Hado Van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Dan Horgan, Bilal Piot, Mohammad Azar, and David Silver. Rainbow: Combining improvements in deep reinforcement learning. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2018.
- [19] Steven Kapturowski, Georg Ostrovski, John Quan, Remi Munos, and Will Dabney. Recurrent experience replay in distributed reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2018.
- [20] Tadashi Kozuno, Eiji Uchibe, and Kenji Doya. Theoretical Analysis of Efficiency and Robustness of Softmax and Gap-Increasing Operators in Reinforcement Learning. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019.
- [21] Clare Lyle, Marc G Bellemare, and Pablo Samuel Castro. A comparative analysis of expected and distributional reinforcement learning. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2019.
- [22] Marlos C Machado, Marc G Bellemare, Erik Talvitie, Joel Veness, Matthew Hausknecht, and Michael Bowling. Revisiting the arcade learning environment: Evaluation protocols and open problems for general agents. *Journal of Artificial Intelligence Research*, 61:523–562, 2018.
- [23] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.
- [24] Rudolf Erich Raspe. Baron Munchhausen’s Narrative of his Marvellous Travels and Campaigns in Russia, 1785.
- [25] Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay. In *International Conference on Representation Learning (ICLR)*, 2016.
- [26] Bruno Scherrer, Mohammad Ghavamzadeh, Victor Gabillon, Boris Lesner, and Matthieu Geist. Approximate modified policy iteration and its application to the game of Tetris. *Journal of Machine Learning Research*, 16:1629–1676, 2015.
- [27] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International Conference on Machine Learning (ICML)*, 2015.
- [28] Zhao Song, Ron Parr, and Lawrence Carin. Revisiting the Softmax Bellman Operator: New Benefits and New Perspective. In *International Conference on Machine Learning (ICML)*, 2019.
- [29] Richard S Sutton. Learning to predict by the methods of temporal differences. *Machine learning*, 3(1):9–44, 1988.

- [30] Yoshihisa Tsurumine, Yunduan Cui, Eiji Uchibe, and Takamitsu Matsubara. Deep dynamic policy programming for robot control with raw images. In *International Conference on Intelligent Robots and Systems (IROS)*, 2017.
- [31] Harm Van Seijen, Mehdi Fatemi, and Arash Tavakoli. Using a logarithmic mapping to enable lower discount factors in reinforcement learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 14134–14144, 2019.
- [32] Nino Vieillard, Tadashi Kozuno, Bruno Scherrer, Olivier Pietquin, Rémi Munos, and Matthieu Geist. Leverage the Average: an Analysis of Regularization in RL. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [33] Nino Vieillard, Bruno Scherrer, Olivier Pietquin, and Matthieu Geist. Momentum in Reinforcement Learning. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.
- [34] Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.
- [35] Derek Yang, Li Zhao, Zichuan Lin, Tao Qin, Jiang Bian, and Tie-Yan Liu. Fully parameterized quantile function for distributional reinforcement learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 6190–6199, 2019.
- [36] Brian D Ziebart. *Modeling Purposeful Adaptive Behavior with the Principle of Maximum Causal Entropy*. PhD thesis, University of Washington, 2010.

Content. These appendices provide the following additional material:

- Appx. A details the derivations made in the paper and proves the stated results.
- Appx. B provides additional experimental details, such as a detailed description of the Munchausen agents, and additional results and visualisations.

Code. All the code used for the experiments is available online at https://github.com/google-research/google-research/tree/master/munchausen_rl.

A Detailed derivation and proofs

This appendix provides additional details regarding the derivation sketched in the main paper as well as the proofs of the stated results:

- Appx. A.1 details the derivation of Soft-DQN.
- Appx. A.2 proves the result that relates Munchausen VI to Mirror Descent VI.
- Appx. A.3 provides and proves component-wise bounds for Munchausen VI, that also apply to Munchausen-DQN.
- Appx. A.4 details the derivation that allows linking the proposed Munchausen approach to the literature.
- Appx. A.5 proves the result quantifying the increase of the action-gap.

First, we recall the notations introduced in the main paper as well as some useful facts about (regularized) MDPs.

We write Δ_X the simplex over the finite set X and Y^X the set of applications from X to the set Y . An MDP is a tuple $\{\mathcal{S}, \mathcal{A}, P, r, \gamma\}$, with \mathcal{S} and \mathcal{A} the state and action spaces (here assumed finite), $P \in \Delta_{\mathcal{S}^{\times \mathcal{A}}}$ the Markovian transition kernel, $r \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ the reward function, uniformly bounded by r_{\max} , and $\gamma \in (0, 1)$ the discount factor. A policy $\pi \in \Delta_{\mathcal{A}}^{\mathcal{S}}$ associates to each state a distribution over actions (a deterministic policy being a special case), and the quality of a policy is quantified by the expected discounted cumulative return, formalized as the state-action value function, $q_{\pi}(s, a) = \mathbb{E}_{\pi}[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) | s_0 = s, a_0 = a]$, the expectation being over trajectories induced by the policy π and the dynamics.

For $f, g \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$, we define a component-wise dot product $\langle f, g \rangle = (\sum_a f(s, a)g(s, a))_s \in \mathbb{R}^{\mathcal{S}}$. This will be used with q -functions and (log-) policies. For $v \in \mathbb{R}^{\mathcal{S}}$, we have $Pv = (\mathbb{E}_{s'|s, a}[v(s')])_{s, a} \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$. We also defined a policy-induced transition kernel P_{π} as $P_{\pi}q = P\langle \pi, q \rangle$. With this, the Bellman evaluation operator is $T_{\pi}q = r + \gamma P_{\pi}q$ and its unique fixed point is q_{π} .

An optimal policy satisfies $\pi_* \in \operatorname{argmax}_{\pi} q_{\pi}$, component-wise, and the associated (unique) optimal value function $q_* = q_{\pi_*}$ satisfies the Bellman equation $q_*(s, a) = r(s, a) + \gamma \mathbb{E}_{s'|s, a}[\max_{a'} q_*(s', a')]$. We write the set of greedy policies as $\mathcal{G}(q) = \operatorname{argmax}_{\pi \in \Delta_{\mathcal{A}}^{\mathcal{S}}} \langle \pi, q \rangle$. We'll also use softmax policies,

$$\pi = \operatorname{sm}(q) \Leftrightarrow \pi(a|s) = \frac{\exp q(s, a)}{\sum_{a'} \exp q(s, a')}.$$

We'll also make use of the entropy of a policy, $\mathcal{H}(\pi) = -\langle \pi, \ln \pi \rangle$, and of the KL between two policies, $\operatorname{KL}(\pi_1 || \pi_2) = \langle \pi_1, \ln \pi_1 - \ln \pi_2 \rangle$. An MDP regularized by a scaled entropy $\tau \mathcal{H}(\pi)$, also known as maximum entropy RL, optimizes for the reward $r - \tau \ln \pi$. It has a unique optimal q -function q_*^{τ} and a unique optimal policy π_*^{τ} , related by $\pi_*^{\tau} = \operatorname{sm}(q_*^{\tau})$; it is related to the solution of the unregularized MDP by $\|q_*^{\tau} - q_*\|_{\infty} \leq \frac{\tau \ln |\mathcal{A}|}{1 - \gamma}$ [15]. We also write q_{π}^{τ} the value function of the policy π in this regularized MDP.

Lastly, by classic properties of the Legendre-Fenchel transform [9, 32], we have $\forall q \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$:

$$\max_{\pi \in \Delta_{\mathcal{A}}^{\mathcal{S}}} \langle q, \pi \rangle + \tau \mathcal{H}(\pi) = \tau \ln \langle 1, \exp \frac{q}{\tau} \rangle = \langle \pi', q \rangle + \tau \mathcal{H}(\pi') \text{ with } \pi' = \operatorname{sm}(q).$$

A.1 Derivation of Soft-DQN

Soft-DQN can be derived from the maximum entropy RL framework. To do so, it is sufficient to follow the derivation that Haarnoja et al. [17] made for SAC. In our case, the actions being discrete, no approximation is necessary for computing the policy (there is no actor), which gives Soft-DQN.

Alternatively, and equivalently, one can derive Soft-DQN as an approximate VI scheme for an MDP regularized by a scaled entropy. The regularized VI scheme is [15, 32]:

$$\begin{cases} \pi_{k+1} = \operatorname{argmax}_{\pi} \langle \pi, q_k \rangle + \tau \mathcal{H}(\pi) \\ q_{k+1} = r + \gamma P \langle \pi_{k+1}, q_k \rangle + \tau \mathcal{H}(\pi_{k+1}) + \epsilon_{k+1} \end{cases}.$$

From Legendre-Fenchel, $\pi_{k+1} = \operatorname{sm}(q_k)$. Using basic calculus, we have

$$\langle \pi_{k+1}, q_k \rangle + \tau \mathcal{H}(\pi_{k+1}) = \langle \pi_{k+1}, q_k \rangle - \tau \langle \pi_{k+1}, \ln \pi_{k+1} \rangle = \langle \pi_{k+1}, q_k - \tau \ln \pi_{k+1} \rangle.$$

Thus, we can write equivalently the regularized VI scheme as

$$q_{k+1} = r + \gamma P \langle \pi_{k+1}, q_k - \tau \ln \pi_{k+1} \rangle + \epsilon_{k+1}, \text{ with } \pi_{k+1} = \operatorname{sm}(q_k),$$

which is basically the Soft-DQN target depicted in Eq. (1).

A.2 Proof of Thm. 1

The proof is similar to the one done in the main paper for the case $\alpha = 1$. Recall Eq. (3), that gives an iteration of M-VI(α, τ):

$$\begin{cases} \pi_{k+1} = \operatorname{argmax}_{\pi \in \Delta_{\mathcal{A}}^S} \langle \pi, q_k \rangle + \tau \mathcal{H}(\pi) \\ q_{k+1} = r + \alpha \tau \ln \pi_{k+1} + \gamma P \langle \pi_{k+1}, q_k - \tau \ln \pi_{k+1} \rangle + \epsilon_{k+1}. \end{cases}$$

Define for any $k \geq 0$ the term q'_k as

$$q'_k = q_k - \alpha \tau \ln \pi_k.$$

By basic calculus, we can rewrite the evaluation step as follows:

$$\begin{aligned} q_{k+1} &= r + \alpha \tau \ln \pi_{k+1} + \gamma P \langle \pi_{k+1}, q_k - \tau \ln \pi_{k+1} \rangle + \epsilon_{k+1} \\ &= r + \alpha \tau \ln \pi_{k+1} + \gamma P \langle \pi_{k+1}, q_k - \alpha \tau \ln \pi_k + \alpha \tau \ln \pi_k - \tau \ln \pi_{k+1} \rangle + \epsilon_{k+1} \\ \Leftrightarrow q'_{k+1} &= r + \gamma P \langle \pi_{k+1}, q'_k - \alpha \tau \ln \frac{\pi_{k+1}}{\pi_k} - (1 - \alpha) \tau \ln \pi_{k+1} \rangle + \epsilon_{k+1} \\ &= r + \gamma P (\langle \pi_{k+1}, q'_k \rangle - \alpha \tau \operatorname{KL}(\pi_{k+1} || \pi_k) + (1 - \alpha) \tau \mathcal{H}(\pi_{k+1})) + \epsilon_{k+1}. \end{aligned}$$

For the greedy step, we have:

$$\begin{aligned} \langle \pi, q_k \rangle + \tau \mathcal{H}(\pi) &= \langle \pi, q_k - \tau \ln \pi \rangle \\ &= \langle \pi, q'_k + \alpha \tau \ln \pi_k - \tau \ln \pi \rangle \\ &= \langle \pi, q'_k - \alpha \tau \ln \frac{\pi}{\pi_k} - (1 - \alpha) \tau \ln \pi \rangle \\ &= \langle \pi, q'_k \rangle - \alpha \tau \operatorname{KL}(\pi || \pi_k) + (1 - \alpha) \tau \mathcal{H}(\pi). \end{aligned}$$

Therefore, we have shown that

$$\begin{aligned} &\begin{cases} \pi_{k+1} = \operatorname{argmax}_{\pi \in \Delta_{\mathcal{A}}^S} \langle \pi, q_k \rangle + \tau \mathcal{H}(\pi) \\ q_{k+1} = r + \alpha \tau \ln \pi_{k+1} + \gamma P \langle \pi_{k+1}, q_k - \tau \ln \pi_{k+1} \rangle + \epsilon_{k+1} \end{cases} \\ &\quad \Downarrow \\ &\begin{cases} \pi_{k+1} = \operatorname{argmax}_{\pi \in \Delta_{\mathcal{A}}^S} \langle \pi, q'_k \rangle - \alpha \tau \operatorname{KL}(\pi || \pi_k) + (1 - \alpha) \tau \mathcal{H}(\pi) \\ q'_{k+1} = r + \gamma P (\langle \pi_{k+1}, q'_k \rangle - \alpha \tau \operatorname{KL}(\pi_{k+1} || \pi_k) - (1 - \alpha) \tau \mathcal{H}(\pi_{k+1})) + \epsilon_{k+1} \end{cases} \end{aligned}$$

This is exactly the update rule of MD-VI($\alpha \tau, (1 - \alpha) \tau$) by Vieillard et al. [32]. Initialized with the same policy π_0 and such that $q'_0 = q_0 - \tau \ln \pi_0$, both algorithms will produce the same sequence of policies (for the same sequence of errors). This is enough for [32, Thm. 1] to apply to M-VI(1, τ), producing the same sequence of policies that MD-VI($\tau, 0$), the result bounding component-wise $q_* - q_{\pi_k}$ (it only involves the computed policy). This is also enough for [32, Thm. 2] to apply to M-VI(α, τ), producing the same sequence of policies that MD-VI($\alpha \tau, (1 - \alpha) \tau$), the result bounding component-wise $q_*^{(1-\alpha)\tau} - q_{\pi_k}$.

A.3 Component-wise bounds for Munchausen VI

We state the component-wise bounds for M-VI, announced in Sec. 3. We recall that they apply to M-DQN, as explained in Sec. 3 (by defining to what corresponds q_k and ϵ_k for M-DQN). First, we provide a bound for the case $\alpha = 1$.

Corollary 1. *Let $(q_k, \pi_k)_{k \geq 0}$ be the sequence of q -functions and policies produced by M-VI($1, \tau$), with π_0 the uniform policy and q_0 such that $\|q_0 - \tau \ln \pi_0\|_\infty \leq \frac{r_{\max}}{1-\gamma}$. Define*

$$E_k = - \sum_{j=1}^k \epsilon_j,$$

$$\text{and } A_k^1 = (I - \gamma P_{\pi_*})^{-1} - (I - \gamma P_{\pi_k})^{-1}.$$

Assume that $\|q_k - \tau \ln \pi_k\|_\infty \leq \frac{r_{\max}}{1-\gamma}$. We have that:

$$0 \leq q_* - q_{\pi_k} \leq \left| A_k^1 \frac{E_k}{k} \right| + \frac{4}{(1-\gamma)^2} \frac{r_{\max} + \tau \ln |\mathcal{A}|}{k} \mathbf{1},$$

with $\mathbf{1} \in \mathbb{R}^{S \times \mathcal{A}}$ the vector whose all components are equal to 1.

Proof. Thanks to Thm. 1, M-VI($1, \tau$) produces the same sequence of policies that MD-VI(λ', τ') with $\lambda' = \tau$ and $\tau' = 0$, and a sequence of q -functions related by $q'_k = q_k - \tau \ln \pi_k$ (q'_k being the q -functions computed by MD-VI(λ', τ')). Thm. 1 of Vieillard et al. [32] thus readily applies, the assumption $\|q'_k\|_\infty \leq \frac{r_{\max}}{1-\gamma}$ translating into $\|q_k - \tau \ln \pi_k\|_\infty \leq \frac{r_{\max}}{1-\gamma}$. \square

Notice that that the assumption that $\|q_k - \tau \ln \pi_k\|_\infty \leq \frac{r_{\max}}{1-\gamma}$ is not strong, it can be ensured by clipping the q_k -values (see also [32, Rk. 1]). Without this, a similar bound would still hold, but with a quadratic dependency of the error term to the horizon, instead of a linear one. Notice that the bound in supremum norm provided in Sec. 3 is a direct corollary of Cor. 1.

Next, we provide a bound for the case $\alpha < 1$.

Corollary 2. *Let $(q_k, \pi_k)_{k \geq 0}$ be the sequence of q -functions and policies produced by M-VI(α, τ), with π_0 the uniform policy, and with $0 \leq \alpha < 1$. For the sequence of policies π_0, \dots, π_k , we define*

$$P_{k:j} = \begin{cases} P_{\pi_k} P_{\pi_{k-1}} \dots P_{\pi_j} & \text{if } j \leq k, \\ I & \text{else,} \end{cases}$$

with $I \in \mathbb{R}^{(S \times \mathcal{A}) \times (S \times \mathcal{A})}$ the identity matrix. We also define

$$A_{k:j}^2 = P_{\pi_*}^{k-j} + (I - \gamma P_{\pi_{k+1}})^{-1} P_{k:j+1} (I - \gamma P_{\pi_j}), \text{ and } E_k^\alpha = (1 - \alpha) \sum_{j=1}^k \alpha^{k-j} \epsilon_j.$$

With these notations, we have

$$0 \leq q_*^{(1-\alpha)\tau} - q_{\pi_{k+1}}^{(1-\alpha)\tau} \leq \sum_{j=1}^k \gamma^{k-j} |A_{k:j}^2 E_j^\alpha| + \gamma^k \left(1 + \frac{1-\alpha}{1-\gamma}\right) \sum_{j=0}^k \left(\frac{\alpha}{\gamma}\right)^j \frac{r_{\max} + (1-\alpha)\tau \ln |\mathcal{A}|}{1-\gamma} \mathbf{1}.$$

Proof. Thanks to Thm. 1, M-VI(α, τ) produces the same sequence of policies that MD-VI(λ', τ') with $\lambda' = \alpha\tau$ and $\tau' = (1-\alpha)\tau$, and a sequence of q -functions related by $q'_k = q_k - \alpha\tau \ln \pi_k$ (q'_k being the q -functions computed by MD-VI(λ', τ')). Thm. 2 of Vieillard et al. [32] thus readily applies, with

$$\beta = \frac{\lambda'}{\lambda' + \tau'} = \frac{\alpha\tau}{\alpha\tau + (1-\alpha)\tau} = \alpha,$$

which gives the stated result. \square

We refer to [32, Sec. 4.2] for an extensive discussion of this bound, but we highlight the fact that it still shows a compensation of errors (through a moving average instead of the average of Cor. 1), something that is desirable.

A.4 Details on related works

First, we relate M-VI to CVI. Recall Eq. (3):

$$\begin{cases} \pi_{k+1} = \operatorname{argmax}_{\pi \in \Delta_{\mathcal{A}}^S} \langle \pi, q_k \rangle + \tau \mathcal{H}(\pi) \\ q_{k+1} = r + \alpha \tau \ln \pi_{k+1} + \gamma P \langle \pi_{k+1}, q_k - \tau \ln \pi_{k+1} \rangle + \epsilon_{k+1}. \end{cases}$$

From the Legendre-Fenchel transform, we have that

$$\pi_{k+1} = \operatorname{sm}\left(\frac{q_k}{\tau}\right) = \frac{\exp \frac{q_k}{\tau}}{\langle 1, \exp \frac{q_k}{\tau} \rangle} \Leftrightarrow \tau \ln \pi_{k+1} = q_k - \tau \ln \langle 1, \exp \frac{q_k}{\tau} \rangle.$$

Injecting this into the evaluation step, we obtain

$$\begin{aligned} q_{k+1} &= r + \alpha \tau \ln \pi_{k+1} + \gamma P \langle \pi_{k+1}, q_k - \tau \ln \pi_{k+1} \rangle + \epsilon_{k+1} \\ &= r + \alpha (q_k - \tau \ln \langle 1, \exp \frac{q_k}{\tau} \rangle) + \gamma P \langle \pi_{k+1}, q_k - (q_k - \tau \ln \langle 1, \exp \frac{q_k}{\tau} \rangle) \rangle + \epsilon_{k+1} \\ &= r + \gamma P(\tau \ln \langle 1, \exp \frac{q_k}{\tau} \rangle) + \alpha (q_k - \tau \ln \langle 1, \exp \frac{q_k}{\tau} \rangle) + \epsilon_{k+1}, \end{aligned}$$

which is exactly Eq. (5), that is a CVI-like update.

It is a classic result that the sum-log-exp tends towards the hard maximum as the temperature goes to zero (this can be also derived from properties of the Legendre-Fenchel transform):

$$\lim_{\tau \rightarrow 0} \tau \ln \sum_a \exp \frac{q_k(s, a)}{\tau} = \max_a q_k(s, a).$$

Using this, the limit of the previous CVI-like update is

$$q_{k+1} = r + \gamma P \langle \pi_{k+1}, q_k \rangle + \alpha (q_k - \langle \pi_{k+1}, q_k \rangle) + \epsilon_{k+1} \text{ with } \pi_{k+1} \in \mathcal{G}(q_k),$$

where we have used that $\max_a q_k(\cdot, a) = \langle \pi_{k+1}, q_k \rangle$ with $\pi_{k+1} \in \mathcal{G}(q_k)$. This is exactly Eq. (6).

A.5 Proof of Thm. 2

This is indeed a corollary of Thm. 1. First, we handle the case $\alpha < 1$. From Thm. 1, we know that M-VI(α, τ) produces the same sequence of policies that MD-VI($\alpha \tau, (1-\alpha)\tau$). From [32, Thm. 2], we now that without error $q'_k = q_k - \alpha \tau \ln \pi_k$ (recall that q'_k is the sequence of q -functions computed by MD-VI) converges to $q_*^{(1-\alpha)\tau}$ and that π_k converges to $\pi_*^{(1-\alpha)\tau}$ (recall that both algorithms produce the same sequence of policies). From this, we can deduce the limit of q_k , the sequence of q -function produced by Munchausen VI:

$$\lim_{k \rightarrow \infty} q_k = q_*^{(1-\alpha)\tau} + \alpha \tau \ln \pi_*^{(1-\alpha)\tau}.$$

From basic properties of regularized MDPs [15], we know that

$$\pi_*^{(1-\alpha)\tau} = \operatorname{sm}\left(\frac{q_*^{(1-\alpha)\tau}}{(1-\alpha)\tau}\right) \Leftrightarrow \ln \pi_*^{(1-\alpha)\tau} = \frac{q_*^{(1-\alpha)\tau}}{(1-\alpha)\tau} - \ln \langle 1, \exp \frac{q_*^{(1-\alpha)\tau}}{(1-\alpha)\tau} \rangle.$$

Therefore, we have that

$$\begin{aligned} \lim_{k \rightarrow \infty} q_k &= q_*^{(1-\alpha)\tau} + \alpha \tau \ln \pi_*^{(1-\alpha)\tau} \\ &= q_*^{(1-\alpha)\tau} + \alpha \tau \left(\frac{q_*^{(1-\alpha)\tau}}{(1-\alpha)\tau} - \ln \langle 1, \exp \frac{q_*^{(1-\alpha)\tau}}{(1-\alpha)\tau} \rangle \right) \\ &= \frac{1+\alpha}{1-\alpha} q_*^{(1-\alpha)\tau} - \frac{\alpha \tau}{1-\alpha} \ln \langle 1, \exp \frac{q_*^{(1-\alpha)\tau}}{(1-\alpha)\tau} \rangle. \end{aligned}$$

Noticing that the log-sum-exp does not depend on the actions, we obtain the stated result.

Next, we handle the case $\alpha = 1$. From Thm. 1, we know that M-VI($1, \tau$) produces the same sequence of policies that MD-VI($\tau, 0$). From [32, Thm. 1], we now that without error $q'_k = q_k - \alpha \tau \ln \pi_k$

converges to q_* and that π_k converges to π_* , the solutions of the unregularized MDP. To simplify and without much loss of generality, assume that this MDP admits a unique optimal policy. As $q_k = q'_k + \alpha \ln \pi_k$, taking the limit we get for any $s \in \mathcal{S}$

$$\lim_{k \rightarrow \infty} q_k(s, a) = \begin{cases} q_*(s, a) & \text{if } \pi_*(a|s) = 1 \\ -\infty & \text{else} \end{cases}.$$

With the adopted convention, this proves the result for the case $\alpha = 1$.

B Additional experimental details and results

This appendix provides a complete description of the Munchausen agents, it gives additional experimental details, and it proposes additional results and visualisations:

- Appx. B.1 provides a complete description of the Munchausen agents, as well as some additional details for the considered metrics (such as human scores for games not reported in the literature) and for the learning setting.
- Appx. B.2 discusses the difference between playing ε -greedy and stochastic policies for Munchausen DQN.
- Appx. B.3 discusses the difference between using 1-step or 3-steps returns in M-IQN.
- Appx. B.4 provides elements of comparison with the original ALE setting.
- Appx. B.5 provides complementary results for the ablation study.
- Appx. B.6 provides complementary comparison results.

B.1 Detailed description of the Munchausen agents

All the agents follow a similar learning procedure, described as a pseudo-code in Alg. 1 for M-DQN. What changes is the loss that is optimized.

M-DQN. Here, we recall the basic workings of M-DQN. It estimates a q -value through an online q -network q_θ of weights θ . Every C steps, the weights are copied to a *target* network $q_{\bar{\theta}}$ of weights $\bar{\theta}$. Transitions (s_t, a_t, r_t, s_{t+1}) are stored in fixed-sized FIFO replay buffer. To collect them, M-DQN interacts with the environment using the policy $\mathcal{G}_\varepsilon(\theta)$, the policy that is ε -greedy with respect to q_θ . M-DQN uses (as DQN) a decay on ε to favour exploration in the beginning of the learning. Each F steps, M-DQN samples a random batch B of transitions from \mathcal{B} and minimizes the following loss, based on the regression target of Eq. (2):

$$\mathcal{L}_{\text{m-dqn}}(\theta) = \hat{\mathbb{E}}_B \left[h \left(r_t + \alpha [\tau \ln \pi_{\bar{\theta}}(a_t|s_t)]_{l_0}^0 + \gamma \sum_{a \in \mathcal{A}} \pi_{\bar{\theta}}(a|s_{t+1}) (q_{\bar{\theta}}(s_{t+1}, a) - \tau \ln \pi_{\bar{\theta}}(a|s_{t+1})) - q_\theta(s_t, a_t) \right) \right], \quad (7)$$

with $\pi_{\bar{\theta}} = \text{sm}(\frac{q_{\bar{\theta}}}{\tau})$ and h the Huber loss function, with a parameter x_h , $h(x) = x^2$ if $x < x_h$ else $|x|$. A pseudo-code detailing the learning procedure is given in Alg. 1.

AL. We have shown in Sec. 3 that AL can be seen as a limiting case of M-DQN, in the limit $\tau \rightarrow 0$. Yet, it cannot be obtained simply by setting $\tau = 0$ in Alg. 1. Instead, we rewrite the minimized loss, according to Sec. 3. Each F steps, AL samples a random batch B of transitions from \mathcal{B} and minimizes the loss

$$\mathcal{L}_{\text{al}}(\theta) = \hat{\mathbb{E}}_B \left[h \left(r_t + \alpha \left(q_{\bar{\theta}}(s_t, a_t) - \max_{a \in \mathcal{A}} q_{\bar{\theta}}(s_t, a) \right) + \max_{a \in \mathcal{A}} q_{\bar{\theta}}(s_{t+1}, a) - q_\theta(s_t, a_t) \right) \right].$$

M-IQN. IQN is a distributional method. It does not estimate directly a q -function, but the distribution of the discounted cumulative rewards, a so-called z -function. Precisely, the z -function $z_\pi \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ of a policy π is a random quantity defined, for each $s, a \in \mathcal{S} \times \mathcal{A}$ as:

$$z_\pi(s, a) = \sum_{t=0}^{\infty} \gamma^t r(s_t, a_t), \text{ with } a_t \sim \pi(\cdot|s_t) \text{ and } s_{t+1} \sim P(\cdot|s_t, a_t) \text{ for } s_0 = s \text{ and } a_0 = a.$$

Algorithm 1 Munchausen DQN

Require: $T \in \mathbb{N}^*$ the number of environment steps, $C \in \mathbb{N}^*$ the update period, $F \in \mathbb{N}^*$ the interaction period.
Initialize θ at random
 $\mathcal{B} = \{\}$
 $\bar{\theta} = \theta$
for $t = 1$ **to** T **do**
 Collect a transition $b = (s_t, a_t, r_t, s_{t+1})$ from $\mathcal{G}_e(\theta)$
 $\mathcal{B} \leftarrow \mathcal{B} \cup \{b\}$
 if $t \bmod F == 0$ **then**
 On a random batch of transitions $B_t \subset \mathcal{B}$, update θ with one step of SGD on $\mathcal{L}_{\text{m-dqn}}$, see (7)
 end if
 if $k \bmod C == 0$ **then**
 $\bar{\theta} \leftarrow \theta$
 end if
end for
return $\mathcal{G}_0(\bar{\theta})$

The q -function can be directly related to it with

$$q_\pi(s, a) = \mathbb{E} [z_\pi(s, a)].$$

A remarkable result is that z_π satisfies a Bellman equation, similarly to q_π , and thus can be estimated with TD. Here, we give a quick overview of IQN, and explain how we modified it. We refer to Dabney et al. [11] for an exact derivation and more details of the original algorithm. IQN estimates the quantile function of z at $\sigma \in [0, 1]$, denoted z_σ . The estimated q -value is then $\tilde{q}(s, a) = \mathbb{E}_{\sigma \sim U_{[0,1]}} [z_\sigma(s, a)]$, this expectation being practically approximated by Monte Carlo. The TD error of IQN at step t , defined with $\sigma, \sigma' \sim U_{[0,1]}$, is:

$$\text{TD}_{\text{IQN}} = r_t + \gamma z_{\sigma'}(s_{t+1}, \pi(s_{t+1})) - z_\sigma(s_t, a_t), \text{ with } \pi(s) = \underset{a \in \mathcal{A}}{\text{argmax}} \tilde{q}(s, a).$$

In practice, $z_{\sigma'}$ is given by a target network, and z_σ by an online network, to be optimized. The loss is then estimated as the empirical mean of the TD errors, by sampling σ and σ' uniformly in $[0, 1]$. In M-IQN, we use an additional Munchausen term in TD error,

$$\text{TD}_{\text{M-IQN}} = r_t + \alpha [\tau \ln \pi(a_t | s_t)]_{l_0}^0 + \gamma \sum_{a \in \mathcal{A}} \pi(a | s_{t+1}) (z_{\sigma'}(s_{t+1}, a) - \tau \ln \pi(a | s_{t+1})) - z_\sigma(s_t, a_t)$$

with $\pi(\cdot | s) = \text{sm}(\frac{\tilde{q}(s, \cdot)}{\tau})$ (that is, the policy is softmax with \tilde{q} , the quantity with respect to which the original policy of IQN is greedy). We use the same parametrization for z as Dabney et al. [11], and all their provided hyperparameters, as implemented in Dopamine. We used the ‘‘Munchausen-RL parameters’’ from Table 2.

Custom log-sum-exp trick. Eq. 7 relies on computing a log-policy, so in our case the log-softmax of a q -values. Such computations are usually done using the ‘‘log-sum-exp trick’’, that allows for numerically stable operations by factorizing a maximum. This trick is widely used in software libraries, for example in TensorFlow [1], used to implement the experiments of this work. With this approach, we use the fact that

$$\tau \ln \pi_{k+1} = q_k - \tau \ln \langle 1, \exp \frac{q_k}{\tau} \rangle,$$

that can be unstable if τ is small. Thus, we compute the log-policy terms using a log-sum-exp-trick as

$$\tau \ln \pi_{k+1} = q_k - v_k - \tau \ln \langle 1, \exp \frac{q_k - v_k}{\tau} \rangle,$$

where we defined $v_k \in \mathbb{R}^S$ as $v_k(s) = \max_a q_k(s, a)$. This is more stable than the one implemented by default, because it takes into account the temperature coefficient.

Parameters. We provide the hyperparameters used in our algorithms in Table 2. We denote neural networks structures as follow: $\text{Conv}_{a,b}^d$ c is a 2D convolutional layer with c filters of size $a \times b$ and of stride d , and $\text{FC } n$ is a fully convolutional layer with n neurons. The parameters of the baseline agents are those reported in Dopamine (with the slight modification of considering 1-step returns instead of n -step returns for IQN, to match the original paper and the algorithm we modify).

Table 2: Parameters used for Munchausen RL agents.

Parameter	Value
Base (Adam) DQN parameters	
C (update period)	8000
F (interaction period)	4
γ (discount)	0.99
$ \mathcal{B} $ (replay buffer size)	10^6
$ B_t $ (batch size)	32
e_t (random actions rate)	0.01 (with a linear decay of period $2.5 \cdot 10^5$ steps)
Q -network structure	$\text{Conv}_{8,8}^4 32 - \text{Conv}_{4,4}^2 64 - \text{Conv}_{3,3}^1 64 - \text{FC } 512 - \text{FC } n_A$
activations	Relu
optimizer	Adam ($lr = 5e - 5$)
Munchausen-RL specific parameters	
τ (entropy temperature)	0.03
α (Munchausen scaling term)	0.9
l_0 (clipping value)	-1
AL specific parameters	
α (advantage scaling term)	0.9

Environment details. We follow the procedures of Machado et al. [22] to train on the ALE. Notably, we perform one training step (a gradient descent step) every 4 frames encountered in the environment. The state of an agent is the concatenation of the last 4 frames, sub-sampled to a shape of (84, 84), in gray levels. We refer to Machado et al. [22] for details on the preprocessing.

Metrics. Here, we recall the definitions of the metrics used to compare algorithms. As an aggregating metric, we use the baseline-normalized score. Every $1M$ frames, we compute the undiscounted return averaged over the last 100 episodes a_k , then we normalized it by a random score r and a baseline score b (score after training for 200M steps). The normalized score is then $\frac{a_k - r}{|b - r|}$. We also use human-normalized scores, when we replace the baseline score by the score of a human. We used human scores reported by [23]. For AirRaid, Carnival, ElevatorAction, JourneyEscape, and Pooyan, not considered in Mnih et al. [23], we averaged scores from game-play posted online by players. For a game-per-game metric, we compute the normalized improvement according to a baseline. The “final score” of an agent is defined as the score averaged over the last 5M frames. The normalized improvement of a final score a w.r.t. the final score of a baseline b is $\frac{a - b}{|b - r|}$. The maximum scores reported in Table 3 are the maximum scores over training, averaged over 100 episodes, averaged over 3 random seeds, obtained during training.

B.2 Comparison of greedy and stochastic policies

Although M-DQN naturally produces stochastic policies, we used the ϵ -greedy one (with respect to q_θ), as explained in Sec. 4. This is motivated by the behaviour of some games. In some games, a random policy fails to gather rewards (as for example Venture or Enduro). The Q -network is initialized with small Q -values, close to zero. Even with the small temperature $\tau = 0$ we consider, the resulting softmax policy is very close to uniform, and the M-DQN fails to collect rewards, and thus receives no signal to learn. On the converse, an ϵ -greedy exploration will have a more (randomly) structured exploration, as the scale of Q -values does not matter in this case. It then succeed to gather rewards, and to learn something. This is exemplified in Fig. 5, left, for the game Enduro.




Figure 5: Comparison of M-DQN with a greedy (blue) or stochastic (orange) interaction policy. **Left:** Enduro. **Right:** Seaquest. On Enduro, the stochastic policy is not able to see any reward signal in the beginning, and learns nothing. On Seaquest, we see that it improves over the greedy policy.




Figure 6: Human-normalized scores of M-DQN greedy and stochastic, mean (**left**) and median (**right**).

On the converse, if the agent manage to get rewards, the M-DQN agent with a stochastic policy will perform more exploration, and a directed one, as it will chose more often actions with high Q -values, thanks to the softmax policy. Consequently, thanks to this less random exploration, it could perform better. We hypothesize that it is what happens for the game Seaquest, shown in Fig. 5, right.

In Fig. 6, we provide the Human-normalized scores of both options, playing with an ϵ -greedy policy or with the more natural stochastic one. We observe that the stochastic policy is slightly better in median. Yet, it improves less games too, and we kept the ϵ -greedy policy for the core results. Improving the stochastic policy, maybe with an adaptive temperature or an adaptive α parameter, is an interesting future direction of research.

B.3 Comparison of 1-step and 3-steps learning in M-IQN

The results in the papers are computed with a version of M-IQN that uses 3-steps learning, and compared to version of IQN that also uses 3-steps learning (as it is by default in the Dopamine library). For completeness, we evaluate M-IQN with 1-steps returns, and compare it to IQN with 1-step returns. The human-normalized scores for these algorithms are reported in Fig.7. Theses results show that (1) n -step learning and M-RL combine efficiently, as M-IQN 3-steps clearly outperforms M-IQN 1-step and (2) that M-IQN alone (with only 1-step returns) yields already high performances, and it particular outperforms – although by a tight margin – the Rainbow baseline, that uses 3-steps returns.

B.4 Element of comparison with the original ALE setting

We explained in Sec. 4 the difference between the ALE setting we consider, more modern and more difficult, compared to the ALE setting often considered, for example for the seminal DQN [23] or for Rainbow [18]. The Rainbow baseline we consider [10] is also not exactly the published one: even if the most important features are included, as deemed by Hessel et al. [18], it does not include all features (such as double Q -learning or dueling architecture).




Figure 7: Human normalized scores of M-IQN, IQN, and Rainbow with different n -steps returns, mean (**left**) and median (**right**). M-IQN, IQN, and Rainbow use 3-steps, while the other two use 1-step.

As a (partial) check, we also evaluated our Munchausen agents, M-DQN and M-IQN, as well as the baselines DQN, IQN and Rainbow, in a setting as close as possible to the one used for the baselines’ publications. Notably, here we did not use sticky actions, making the environment deterministic, and we end an episode whenever the agent loses a life, instead of when it encounters a game-over. We also use hyperparameters provided in the original publications, the only difference being that we used a target update period of 10000 steps instead of 8000. We did so on the Asterix game, the results being depicted in Fig. 8.

On Fig. 8, left, we can observe DQN and M-DQN. The result for DQN is normal, despite the apparent “crash”, see for example the training curves in [18] (notice also that it is often the best scores over training which is reported, instead of the final one, as in our Tab. 3 or in the seminal DQN publication [23]). We can observe that M-DQN performs much better than DQN, without falling, and that the score is close to the one of M-DQN in the more difficult setting (15k vs 19k in the more difficult setting).

On Fig. 8, right, we can observe Rainbow, IQN and M-IQN. All algorithms perform pretty well. For example, Rainbows reaches roughly 350k, comparable to the original publication⁷. This is much more than in our setting, where Rainbow reaches only 18k, suggesting that the original setting is easier. We can also see that IQN works well (and somehow surprisingly better than in the original publication, compared to Rainbow), and that M-IQN works better than both IQN and Rainbow.

An interesting thing is to see how the methods degrades (roughly) when going from the agent is trained in the considered setting, compared to the original one. Rainbow goes from 350k to 18k (5% of the original scores), IQN goes from 350k to 33k (10%), while M-DQN goes from 15k to 17k (113%) and M-IQN goes from 350k to 50k (17%). This suggests that M-RL might be more stable over environments.

For sure, this discussion only holds for one game, and no general conclusion can be drawn. Yet, it suggests a few things, the ALE setting we consider is more difficult, among other advantages [22], the Rainbow baseline we consider is correct, and M-RL seems to be more stable.

B.5 Additional results on the ablation study

We provide complementary results regarding the ablation study:

- Fig. 9 p. 22 reports the Rainbow-normalized scores of the ablation (instead of the Human-normalized ones in the main paper, Fig. 3).
- Fig. 11 p. 24 shows the normalized improvements of all ablations with respect to DQN.
- Fig. 12 p. 25 reports all learning curves an the 60 Atari games for the ablation.

⁷The setting is still not exactly the same, due to less enhancements in the Dopamine’s Rainbow, a different codebase, but also a difference in the start (human start vs no-op for Rainbow, straight start for us), and possibly a different ROM, which cannot be checked.




Figure 8: Scores of different agent on the game Asetrix, using the original ALE. **left:** M-DQN and DQN. **right:** Rainbow, IQN and M-IQN.




Figure 9: Rainbow-normalized ablation study results. **Left:** mean. **Right:** median.

The Rainbow-normalized scores (Fig. 9) confirms the Human-normalized ones (Fig. 3). The scores themselves are different (due to a different normalization), but the order of the different variations and their gaps is comparable.

Fig. 11 provides a summary of the per-game improvement, while Fig. 12 provides all related learning curves (Fig. 11 summarizing what the results are after 200M frame). We can observe that M-DQN is not always the best performing agent. Yet, it is very often competitive with the best performing ablation (when M-DQN does not perform the best), and the ablation that surpasses M-DQN is highly game-dependent. Overall, M-DQN is consistently the best performing agent over the whole suite of games, as confirmed by Fig. 3 or Fig. 9 both in mean and median Rainbow and Human-normalized scores.

AL performs pretty well (even if less well than M-DQN). Yet, Munchausen-RL is more general, as it consists only in adding a scaled log-policy term to the reward. We’ve shown in the main paper how it can be readily applied to agents that does not even consider stochastic policies. On the converse, ALE relies heavily on being able to compute the maximum Q -value, something which could not be easily extended to continuous actions, contrary to the Munchausen principle. We let this as an interesting direction for future work. In both average and mean (Fig. 3 and 9), Soft-DQN is the worst ablation, despite being much better in a few games (for example, Amidar or Jamesbond). Again, the temperature was not specifically tuned for Soft-DQN, but it is on par with the close literature (see discussion in Sec. 4). This suggests that the maximum entropy RL principle alone might not be sufficient, especially when one observes the significant improvement that the Munchausen term brings to it (or, implicitly, adding KL regularization to the entropy term). We also notice again that Adam DQN works surprisingly well, compared to the original DQN. This is a very interesting finding, and it suggests that Adam DQN should be considered as a better baseline than the seminal DQN.




Figure 10: Rainbow-normalized scores. **Left:** mean. **Right:** median.

B.6 Additional comparison results

For completeness, we provide additional comparison results:

- In addition to the Human-normalized results of Fig. 1, we provide a Rainbow-normalized comparison of the Munchausen agents with respect to DQN, CS1, IQN and Rainbow in Fig. 1.
- In addition to the per-game normalized improvement of a Munchausen agent with respect to its natural baseline (Fig. 4), we provide the per-game improvement for M-DQN over DQN, CS1, IQN and Rainbow in Fig. 13, as well as the per-game improvement of M-IQN over the same baselines in Fig. 14.
- We provide a summary of all best scores (among training, averaged over 3 seeds), for all games on all agents, in Table 3 p. 28. M-IQN obtains the most highest-ranking scores among all the considered baselines (including the human one).
- For completeness, we report all learning curves of the Munchausen agents and the considered baselines, for the full set of Atari games, in Fig. 15.

These additional results confirm the observations made in the main paper.




Figure 11: Per games N.I./DQN of the ablation study.




Figure 12: All averaged training scores of the ablation. M-DQN in blue, AL in orange, Soft-DQN in green, DQN Adam in red, and DQN in dashed purple.

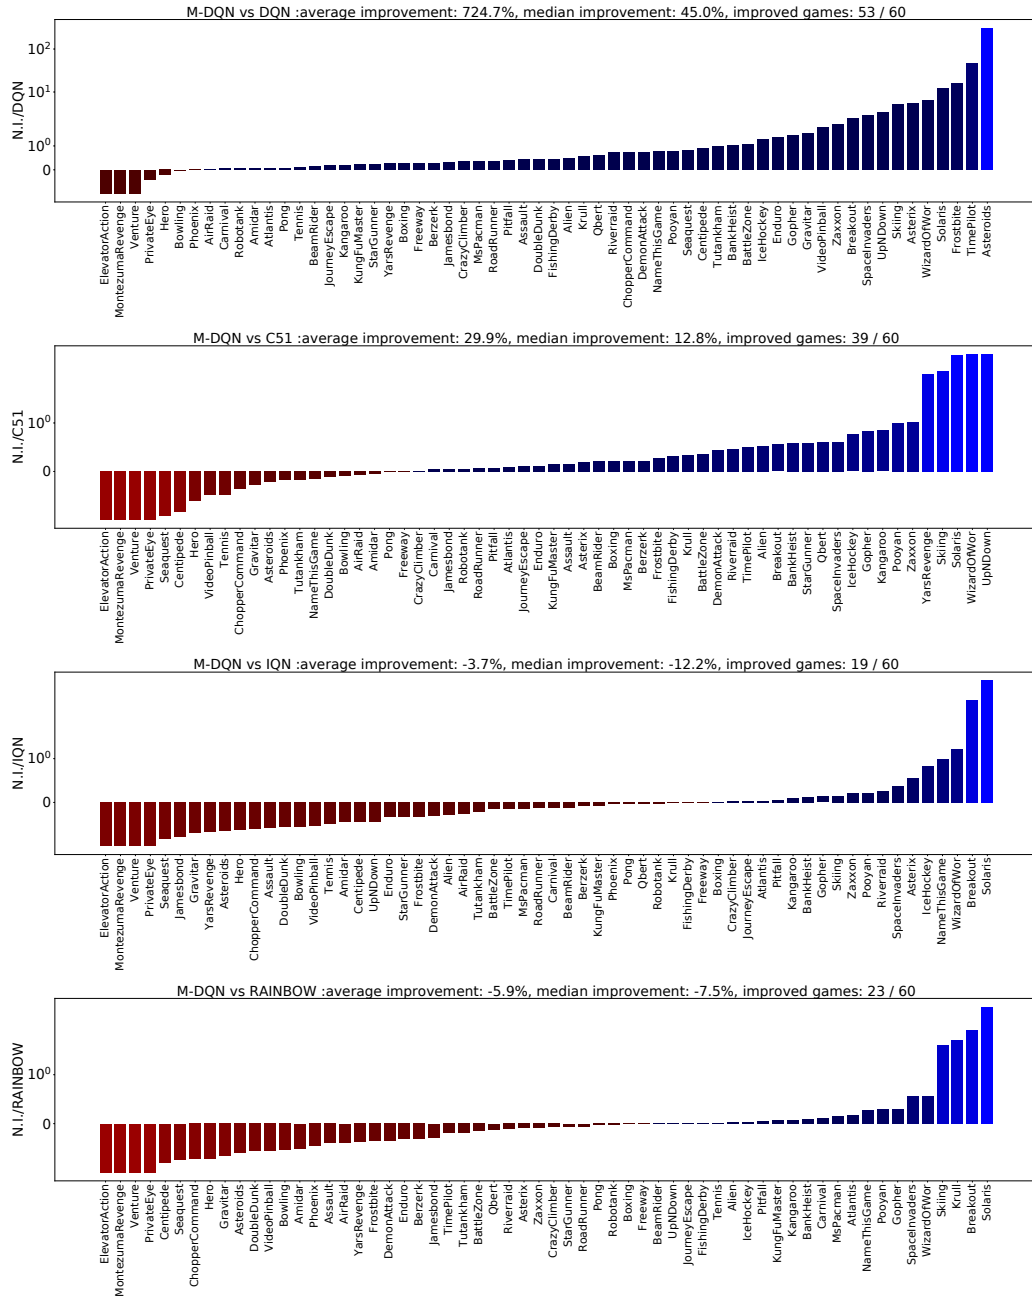


Figure 13: Normalized Improvement of M-DQN vs DQN, C51, IQN, and Rainbow.




Figure 14: Normalized Improvement of M-IQN vs DQN, C51, IQN, and Rainbow.

Table 3: Maximum scores obtained during training (averaged over 100 episodes and 3 random seeds). The bottom line counts the number of games on which an algorithm or a human performs the best.

	random	human	IQN	DQN	RAINBOW	M-DQN	M-IQN
AirRaid	400	3000	15077	7700	14056	8914	19111
Alien	228	7128	5119	2533	3587	3795	4492
Amidar	6	1720	2442	1222	2630	1423	1875
Assault	222	742	4902	1573	3511	2165	7504
Asterix	210	8503	10965	3433	18367	17238	49865
Asteroids	719	47389	1616	828	1489	1150	1685
Atlantis	12850	29028	893764	919622	838590	939533	918183
BankHeist	14	753	1073	704	1148	1190	1292
BattleZone	2360	37188	41475	18667	40895	36509	52517
BeamRider	364	16926	7365	5852	6529	6745	12775
Berzerk	124	2630	662	559	842	608	736
Bowling	23	161	46	33	49	37	32
Boxing	0	12	98	82	99	98	99
Breakout	2	30	159	127	120	331	320
Carnival	380	4000	5712	4860	5069	5022	5588
Centipede	2091	12017	3816	3337	6618	4134	4371
ChopperCommand	811	7388	9301	2852	12844	4507	4573
CrazyClimber	10780	35829	137201	109635	147743	140156	150783
DemonAttack	152	1971	15433	6411	17802	12114	68825
DoubleDunk	-19	-16	21	-6	22	0	22
ElevatorAction	0	3000	67224	1723	79968	4215	89237
Enduro	0	860	2270	815	2230	1643	2332
FishingDerby	-92	-39	45	9	43	44	55
Freeway	0	30	34	26	34	34	34
Frostbite	65	4335	8061	1186	8572	5453	9538
Gopher	258	2412	12108	6044	10641	14728	27469
Gravitar	173	3351	1350	330	1272	550	1134
Hero	1027	30826	36583	17330	46764	13824	26037
IceHockey	-11	1	-0	-6	2	0	12
Jamesbond	29	303	3596	589	1106	814	1637
JourneyEscape	-18000	-1000	-1252	-2668	-959	-938	-806
Kangaroo	52	3035	12872	12192	13460	14067	10939
Krull	1598	2666	8910	6410	6229	8912	10703
KungFuMaster	258	22736	33348	24495	27900	29607	27119
MontezumaRevenge	0	4753	500	2	500	0	0
MsPacman	307	6952	5225	3471	4027	4544	6029
NameThisGame	2292	8049	9129	7348	9229	11807	12761
Phoenix	761	7243	5137	5651	8605	5140	5327
Pitfall	-229	6464	-3	-17	-1	0	0
Pong	-21	15	20	17	20	19	19
Pooyan	500	1000	5339	3535	5640	6396	13096
PrivateEye	25	69571	6852	1004	21532	121	100
Qbert	164	13455	16995	10399	18503	16415	14739
Riverraid	1338	17118	15554	12051	21091	19346	16271
RoadRunner	12	7845	59443	39468	55300	51866	61269
Robotank	2	12	67	61	66	66	73
Seaquest	68	42055	19170	2133	11362	2666	23885
Skiing	-17098	-4337	-11035	-15712	-20518	-9671	-10336
Solaris	1236	12327	2204	1955	2438	5169	5765
SpaceInvaders	148	1669	5452	1850	4420	7504	13871
StarGunner	664	10250	80362	45015	57909	55100	65757
Tennis	-24	-8	23	-0	0	0	0
TimePilot	3568	5229	11887	3768	12283	10590	15155
Tutankham	11	168	256	132	245	200	207
UpNDown	533	11693	74659	10348	39065	45738	216080
Venture	0	1188	1430	52	1579	19	1101
VideoPinball	0	17668	485551	177488	513484	368930	625118
WizardOfWor	564	4756	6208	2597	8201	12517	13644
YarsRevenge	3093	54577	85762	24389	45567	29792	111583
Zaxxon	32	9173	11761	4825	15089	13905	19080
Best	0	14	7	0	8	3	28




Figure 15: All averaged training scores. M-DQN in blue, M-IQN in orange, IQN in dashed green, Rainbow in dashed red, DQN in dashed purple, and C51 in dashed brown.