

PROBING TRANSFER IN DEEP REINFORCEMENT LEARNING WITHOUT TASK ENGINEERING

Andrei A. Rusu, Sebastian Flennerhag, Dushyant Rao, Razvan Pascanu, Raia Hadsell

DeepMind, UK

{andrei, flennerhag, dushyantr, razp, raia}@deepmind.com

ABSTRACT

We evaluate the use of original game curricula supported by the Atari 2600 console as a heterogeneous transfer benchmark for deep reinforcement learning agents. Game designers created curricula using combinations of several discrete modifications to the basic versions of games such as Space Invaders, Breakout and Freeway, making them progressively more challenging for human players. By formally organising these modifications into several factors of variation, we are able to show that Analyses of Variance (ANOVA) are a potent tool for studying the effects of human-relevant domain changes on the learning and transfer performance of a deep reinforcement learning agent. Since no manual task engineering is needed on our part, leveraging the original multi-factorial design avoids the pitfalls of unintentionally biasing the experimental setup. We find that game design factors have a large and statistically significant impact on an agent’s ability to learn, and so do their combinatorial interactions. Furthermore, we show that zero-shot transfer from the basic games to their respective variations is possible, but the variance in performance is also largely explained by interactions between factors. As such, we argue that Atari game curricula offer a challenging benchmark for transfer learning in RL, that can help the community better understand the generalisation capabilities of RL agents along dimensions which meaningfully impact human generalisation performance. As a start, we report that value-function finetuning of regularly trained agents achieves positive transfer in a majority of cases, but significant headroom for algorithmic innovation remains. We conclude with the observation that selective transfer from multiple variants could further improve performance.

1 INTRODUCTION

A key open challenge in artificial intelligence is training reinforcement learning (RL) agents which generally achieve high returns when faced with critical changes to their environments (Schaul et al., 2018), motivated by impressive flexibility of animal and human learning. One way to approach the problem is through the prism of *generalisation* across related but distinct environments, also called *transfer learning* (Pan & Yang, 2009; Taylor & Stone, 2009) and comprehensively reviewed in the RL setting by Zhu et al. (2020). Many purpose-built benchmarks serve investigations into more specific research questions, e.g. transfer learning in particular cases where additional assumptions hold. While useful for progress, the challenge of transfer learning in the more general case remains.

Motivated by visual observation similarity, Machado et al. (2018) suggest using the newest iteration of the Atari Learning Environment (ALE) (Bellemare et al., 2013) to study the transfer learning between single-player game variants, or “flavours”, found in the curricula of many Atari game titles. We will use the terms *default* or *basic* game interchangeably to refer to the environment recommended by respective manuals as the entry point. We call all other distinct games *variations* or *variants* of their respective *default* game. Hence, each Atari game title we consider provides a curriculum consisting of a *default* and its *variants*, all designed to teach and challenge human players in novel ways.

Studying transfer *within curricula* ensures that environments are related, and that meaningful knowledge reuse should be possible and beneficial. Interestingly, differences in game variant dynamics, subtle changes in observations which are crucial for optimal behaviour, novel environment states, as well as new player abilities challenge *unified* approaches to transfer learning across variations. Farebrother et al. (2018) argue that Deep Q-Network (DQN) agents (Mnih et al., 2015), trained with appropriate regularisation, can be effective in zero-shot and finetuning transfer scenarios. In this work we study the learning and transfer performance of an updated version of the DQN agent, called Rainbow-IQN (Toromanoff et al., 2019). We use ANOVA to quantitatively confirm the suspected link between game design factors and agent performance. While current approaches occasionally achieve meaningful transfer from default games, they have limited success for variations with several modifications. Our analyses reveal this is due to strong interaction

effects between factors, not just the isolated effects of the modifications they introduce. This reinforces the case for agents leveraging these systematic curricula, originally designed for human players, through transfer learning.

Contributions: (1) We show that discrete changes to Atari game environments, challenging for human players, also modulate the performance of a popular model-free deep reinforcement learning algorithm starting from scratch. (2) Zero-shot transfer of policies trained on one game variation and tested on others can be significant, but performance is far from uniform across respective environments. (3) Interestingly, zero-shot transfer variance from default game experts is also well explained by game design factors, especially their interactions. (4) We empirically evaluate the performance of value-function finetuning, a general transfer learning technique compatible with model-free deep RL, and confirm that it can lead to positive transfer from basic game experts. (5) We point out that more complex challenges of transfer with deep RL are captured by Atari game variations, e.g. appropriate source task selection, fast policy transfer and evaluation using experts, as well as data efficient behaviour adaptation more generally.

2 BACKGROUND

Reinforcement Learning. A Markov Decision Processes (MDP) (Puterman, 1994) is the classic abstraction used to characterise the sequential interaction loop between an action taking agent and its environment, which responds with observations and rewards (Sutton & Barto, 2018). Formally, a finite MDP is a tuple $\mathcal{M} = \langle \mathcal{X}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \gamma \rangle$, where \mathcal{X} is the state space, \mathcal{A} is the action space, both finite sets, $\mathcal{T} : \mathcal{X} \times \mathcal{A} \times \mathcal{X} \mapsto [0, 1]$ is the stochastic transition function which maps each state and action to a probability distribution over possible future states $\mathcal{T}(x, a, x') = P(x'|x, a)$, $\mathcal{R} : \mathcal{X} \times \mathcal{A} \times \mathcal{X} \mapsto \mathbb{R}$ is the reward distribution function, with $r : \mathcal{X} \times \mathcal{A} \mapsto \mathbb{R}$ the expected immediate reward $r(x, a) = \mathbb{E}_{\mathcal{T}}[\mathcal{R}(x, a, x')]$, and $\gamma \in [0, 1]$ is the discount factor (Bellman, 1957). A stochastic policy function is the action selection strategy $\pi : \mathcal{X} \times \mathcal{A} \mapsto [0, 1]$ which maps states to probability distributions over actions $\pi(x) = P(a|x)$. The discounted sum of future rewards is the random variable $Z_{\mathcal{M}}^{\pi}(x, a) = \sum_{t=0}^{\infty} \gamma^t r(x_t, a_t)$, where $x_0 = x$, $a_0 = a$, $x_t \sim \mathcal{T}(\cdot|x_{t-1}, a_{t-1})$ and $a_t \sim \pi(\cdot|x_t)$. Given an MDP \mathcal{M} and a policy π , the value function is the expectation over the discounted sum of future rewards, also called the expected return: $V_{\mathcal{M}}^{\pi}(x) = \mathbb{E}[Z_{\mathcal{M}}^{\pi}(x, \pi(x))]$. The goal of Reinforcement Learning (RL) is to find a policy $\pi_{\mathcal{M}}^* : \mathcal{X} \mapsto \mathcal{A}$ which is optimal, in the sense that it maximises expected return in \mathcal{M} . The state-action Q-function is defined as $Q_{\mathcal{M}}^{\pi}(x, a) = \mathbb{E}[Z_{\mathcal{M}}^{\pi}(x, a)]$ and it satisfies the Bellman equation $Q_{\mathcal{M}}^{\pi}(x, a) = \mathbb{E}_{\mathcal{T}}[\mathcal{R}(x, a, x')] + \gamma \mathbb{E}_{\mathcal{T}, \pi}[Q_{\mathcal{M}}^{\pi}(x', a')]$ for all states $x \in \mathcal{X}$ and actions $a \in \mathcal{A}$. Mnih et al. (2013) adapted a reinforcement learning algorithm called Q-Learning (Watkins, 1989) to train deep neural networks end-to-end, mastering several Atari 2600 games, with inputs consisting of high-dimensional observations, in the form of console screen pixels, and differences in game scores. For more details please consult Appendix A.

Note that value functions depend critically on all aspects of the MDP. For any policy π , in general $Q_{\mathcal{M}}^{\pi} \neq Q_{\mathcal{M}'}^{\pi}$ for MDPs \mathcal{M} and \mathcal{M}' defined over the same state and action sets, with $\mathcal{T} \neq \mathcal{T}'$ or $\mathcal{R} \neq \mathcal{R}'$. Even if differences in dynamics or rewards are isolated to a subset of $\mathcal{X} \times \mathcal{A}$, changes may be induced across the support of $Q_{\mathcal{M}'}^{\pi}$, since value functions are expectations over sums of discounted future rewards, issued according to \mathcal{R}' , along sequences of states decided entirely by the new environment dynamics \mathcal{T}' when following a fixed behaviour policy π . Nevertheless, many particular cases of interest exist, which we discuss below.

Transfer learning. Described and motivated in its general form by Caruana (1997); Thrun & Pratt (1998); Bengio (2012), the goal of transfer learning is to use knowledge acquired from one or more *source* tasks to improve the learning process, or its outcomes, for one or more *target* tasks, e.g. by using fewer resources compared to learning from scratch. When this is achieved, we call it *positive* transfer. We often further qualify transfer learning by the metric used to measure specific effects. Several transfer learning metrics have been defined for the RL setting (Taylor & Stone, 2009), but none capture all aspects of interest on their own. A first metric we use is “jumpstart” or “zero-shot” transfer, which is performance on the target task before gaining access to its data. Another highly relevant metric is “performance with fixed training epochs” (Zhu et al., 2020), defined as returns achieved in the target task after using fixed computation and data budgets under transfer learning conditions.

One way to classify such approaches in RL is by the format of knowledge being transferred (Zhu et al., 2020), commonly: datasets, predictions, and/or parameters of neural networks encoding representations of observations, policies, value-functions or approximate state-transition “world models” acquired using source tasks. Another way to classify transfer learning approaches in RL is by their respective sets of *assumptions* about source and target tasks, also well illustrated by their associated benchmark domains: (1) Differences are limited to observations, but the underlying MDP is the same, e.g. domain adaptation and randomisation (Tobin et al., 2017), mastered through the generalisation abilities of large models (Cobbe et al., 2019; 2020). (2) MDP states and dynamics are the same, but reward functions are different: successor features and representations (Barreto et al., 2017). (3) Overlapping state spaces with similar dynamics, e.g. multitask training and policy transfer (Rusu et al., 2016a; Schmitt et al., 2018), contextual parametric approaches, e.g. UVFAs (Schaul et al., 2015) and collections of skills/goals/options (Barreto et al., 2019). (4) Sus-

pected but unqualified overlaps between tasks (Parisotto et al., 2016; Rusu et al., 2016b; 2017). (5) Large, curated collections of similar environments designed to facilitate complex transfer and fast adaptation through meta-learning (Yu et al., 2020; Hospedales et al., 2020). All these works capture important sub-problems of interest, and clever design of specialised benchmarks has greatly aided progress. Atari game variations (Machado et al., 2018) offer the exciting prospect of direct comparisons between generalisations of transfer learning methods originally developed under different sets of assumptions, by measuring their performance along dimensions of variation which are meaningful and challenging for human players, one of many interesting criteria cutting across specialised paradigms.

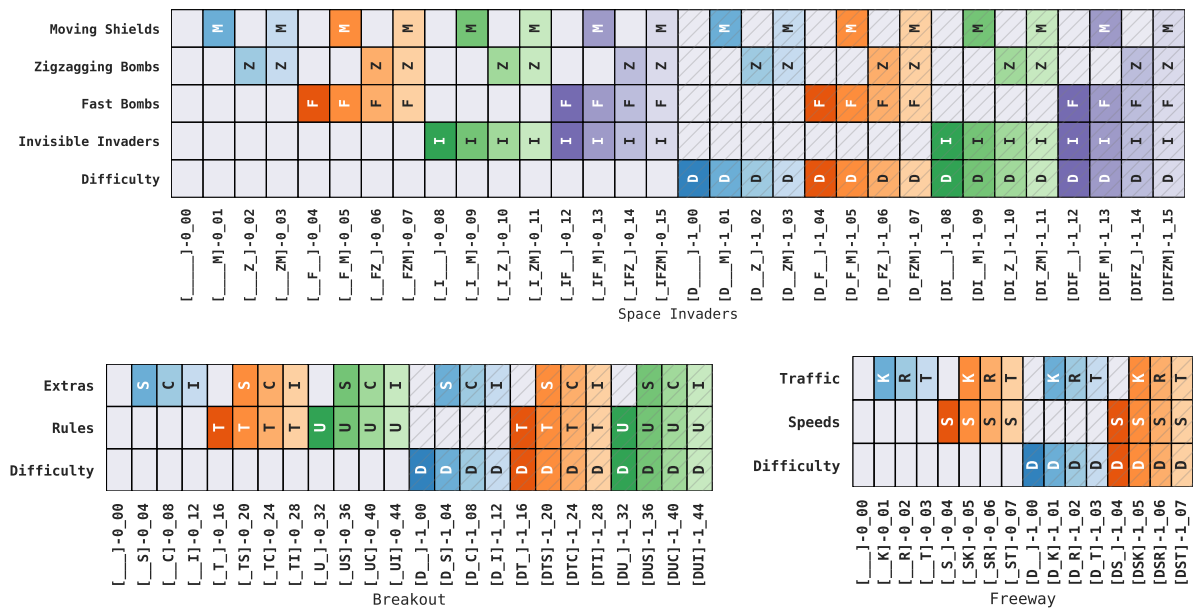


Figure 1: Variant naming convention and factorial design matrices for Space Invaders (top), Breakout (bottom left) and Freeway (bottom right). All factors of variation are categorical and highlighted only when taking different values from default games. Binary factors are abbreviated by their initial, e.g. “Moving Shields” is plotted as ‘M’. Non-binary factors are: Breakout “Rules” with additional values: ‘T’ for “Timed Breakout”, or ‘U’ for “Breakthru”; Breakout “Extras” with additional values: ‘S’ for “Steerable”, ‘C’ for “Catch”, or ‘I’ for “Invisible”; Freeway “Traffic” with levels: ‘K’ for “Thick”, ‘R’ for “Thicker”, or ‘T’ for “Thickest”. Colours indicate the “game mode”, and hatching denotes the “difficulty” switch being activated. Colours and horizontal labels correspond to those used in subsequent figures. Labels include factor abbreviations to supplement the ALE naming conventions.

2.1 ALE ATARI AS A TRANSFER LEARNING BENCHMARK

With the first version of the Arcade Learning Environment (ALE), Bellemare et al. (2013) gave access to over 50 different Atari 2600 game titles through a unified observation and action interface, modelled after the standard reinforcement learning loop. The ALE proved an excellent development test-bed for building general deep RL agents, in no small part due to its diversity and being devoid of experimenter bias, since games were not modified by researchers.

The second and latest version of the ALE (Machado et al., 2018) opens up a wealth of game variations for transfer learning research. This is achieved by emulating functions of “difficulty” and “game select” switches for single-player games. The original cartridges of many game titles came with variations which could be selected and played using these switches. We assign unique identifiers of game variants using the X_YZ notation, with $X \in \{0, 1\}$ denoting the position of the “difficulty” switch, and $YZ \in \{00, 01, 02, \dots\}$ indicating the selected “game mode”. Value ranges are game specific and identical to those used by the ALE code-base (Machado et al., 2018). For example, the entry game version of each title is denoted as 0_00, which we call the default or basic game. All game variations are different environments that still feature the main concepts of the default. Furthermore, the variants were designed to serve as curricula for human players, hence positive transfer should be possible. Most importantly for our purposes, the ALE remains largely free of experimenter bias. Farebrother et al. (2018) selected a few representative variants from four game titles for experiments. We aim to analyse entire curricula, thus we consider the top three most popular titles from their list, according to sales, and use all their variations, for a total of 72 distinct environments.

The original designers created game variations using combinations of discrete modifications to a default “entry” game, qualitatively described in accompanying game manuals. We organise these discrete modification into formal factors of variation, following closely the original game design matrices sometimes explicitly plotted in manuals. In Figure 1 we formalise variant naming conventions relative to default games. We briefly explain the meanings of all design factors below, to help build an intuitive understanding of this heterogeneous collection of transfer learning scenarios.

Space Invaders. The player controls an Earth-based laser cannon which can move horizontally along the bottom of the screen. There is a grid of approaching aliens above the player, shooting laser bombs, and the objective of the game is to eliminate all the aliens before they reach the Earth, especially their Command Ships. Three destructible shields above the cannon offer some protection. The game ends once any alien reaches the Earth or the player’s cannon is hit a third time with laser bombs. Game variants are combinations of five binary factors which, when activated, modify the default: (1) The difficulty switch widens the player’s laser cannon, making it an easier target for enemy laser bombs. (2) Shields move, and thus are harder to use. (3) Enemy laser bombs zigzag, which makes them harder to predict. (4) Enemy laser bombs drop faster. (5) Invaders are invisible, and only hitting one with the laser briefly illuminates the others. This creates a total of 32 variants of Space Invader.

Breakout. In order to achieve the best score in any variant, the players need to completely break walls with six layers of coloured bricks by knocking off said bricks with a ball, served and bounced towards the wall with a controllable paddle, which they can move horizontally along the bottom of the screen. The ball will also bounce off of screen edges, except for the bottom edge, where balls either come in contact with the paddle or are lost. Players have a total of five balls at their disposal, and the game ends when all balls are lost. Points are scored when bricks are knocked off the wall, and the ball accelerates on contact with bricks in the top three rows or after twelve consecutive hits. Game variants are created by all combinations of three factors: (1) The binary difficulty switch reduces the paddle’s width by a quarter, making it easier to miss the ball. (2) The precise rules which determine how points accumulate: (2-a) with standard “Breakout” rules, e.g. in the default game (0_00), the player must completely break down two walls of bricks, one after the other, while loosing the fewest balls; (2-b) under “Timed Breakout” rules, the player must completely break a single wall as fast as possible, no matter how many of the five balls are used; (2-c) with “Breakthru” rules, the player needs to break two walls consecutively, but the ball does not bounce off bricks, unlike in previous variants; the ball keeps going through the wall, quickly picking up speed and accumulating points. (3) How the ball is aimed at the wall and “extras”: (3-a) the ball simply bounces off the paddle at a position dependent angle; (3-b) the player can also steer the ball in flight; (3-c) the player is also able catch the ball and release it at a slower speed; (3-d) the wall is invisible and is only briefly illuminated when a brick is knocked off, so the player needs to remember its configuration in order to aim effectively. All factors together define 24 variants of Breakout.

Freeway. In all variants, the goal of the player is to safely get a chicken across ten lanes of freeway traffic as many times as possible in 2 minutes and 16 seconds. Variants are constructed as combinations of three design factors: (1) The difficulty switch controls whether the chicken is knocked back one lane, or all the way to the kerb, when hit by incoming vehicles. (2) Traffic “thickness”, defined as four levels of traffic density. (3) Vehicle speeds across lanes, either constant or randomised. Hence, there are 16 variants of Freeway.

3 METHODOLOGY

Experimental setup. Following the latest ALE benchmark recommendations (Machado et al., 2018), information about player “lives” is not disclosed to agents, and stochasticity is introduced using randomised action replay with probability 25%, also known as “sticky actions”. Irrespective of redundancies, agents act using the full set of 18 actions. Environment observations (Atari screen frames) were pre-processed according to standard practice (Mnih et al., 2015; Hessel et al., 2018). We report agent returns at the end of training, averaged over the last 2 million steps.

Expert Agent Training. We use the Rainbow-IQN model-free deep reinforcement learning algorithm (Hessel et al., 2018; Dabney et al., 2018; Toromanoff et al., 2019) since it is available to the community in several implementations, e.g. Castro et al. (2018), and is effective with widely available commodity hardware and open-source software.

Rainbow (Hessel et al., 2018) collects a number of improvements to DQN (Mnih et al., 2013; 2015), of which we used Double Q-Learning (Van Hasselt et al., 2016), Prioritised Replay (Schaul et al., 2016) and multi-step learning (Sutton & Barto, 2018). Following Castro et al. (2018), we did not use the dueling network architecture (Wang et al., 2016) or noisy networks (Fortunato et al., 2018). We replaced the distributional RL approach C51 (Bellemare et al., 2017) with its more general form (IQN), since it has been shown to be superior (Dabney et al., 2018; Toromanoff et al., 2019).

We used the standard limit on environment interactions of 200 million steps (Mnih et al., 2015; 2016; Dabney et al., 2018; Machado et al., 2018; Toromanoff et al., 2019). Our study is the first to report Rainbow-IQN results for Atari game variants, hence we use independent hyper-parameter selection for expert training and finetuning experiments.

Expert Agent Finetuning. The experimental setup for finetuning follows closely that for agent training, except for two important differences: (1) We used 10 million environment steps to adapt to new variants, following [Farebrother et al. \(2018\)](#), which is $20\times$ less data than what variant-experts are trained with. The aim of transfer learning is to reduce resources needed for acquiring new knowledge and behaviours, hence we are interested in improving sample complexity using transfer. (2) The hyper-parameter grid was slightly adapted in order to improve chances of fast learning in this reduced data regime. Further details, including parameter grids, are listed in [Appendix B](#).

Statistical Analyses. A multi-factor Analyses of Variance (ANOVA) ([Girden, 1992](#)) is a statistical procedure which investigates the influence that two or more independent variables, or factors, have on a dependent variable. The factors can take two or more categorical values, and experimental designs are often balanced: they use equal numbers of independent observations for all combinations of factor values, no less than three. Other assumptions are normality of deviations from group means, and equal or similar variances, see the discussion in [Appendix C](#). ANOVA can be used to reveal statistically significant evidence against the *null-hypothesis* that group means are all the same. We study the influence of game design factors, discrete modification to default games, on average returns of learned policies.

Study Limitations. The relationship between hyper-parameters and quality of policies learned with deep RL is complex and poorly understood. [Mnih et al. \(2015\)](#) used task-agnostic DQN hyper-parameters, tuned across a few basic games. Later works introduce several interacting modification, including changes to sets of hyper-parameters ([Schaul et al., 2016](#); [Dabney et al., 2018](#); [Toromanoff et al., 2019](#)). With the risk of maximisation bias, we used grid searches per game variant in order to mitigate the greater risks of incorrect inferences about Atari curricula due to poor hyper-parameter settings or divergence. Future works may perform fine-grained sensitivity analyses, using more than two random seeds. Due to computational limitations, we instead report the returns of the top-three agents from our grids.

4 EXPERIMENTS

We would like to characterise the performance of general transfer learning strategies across heterogeneous scenarios, designed to progressively challenge human players. While the shapes of learning curves are expected to be different between game variants, agent performance is thought to be bounded only by its inductive biases, the learning algorithm and practical limitation on resources, such as computation or interactions ([Silver et al., 2021](#)). We say that some game variants is “harder” for a given agent if the average returns of learned policies are lower given equal resource budgets. We aim to explain this in terms of the factors of variation which combinatorially define environments within curricula, knowing that they also impact human player performance. Comparing raw scores across the benchmark is difficult if the underlying tasks are “harder” to learn from scratch, or if variants have different scoring scales, as is sometimes the case here. Hence, we must first establish the performance of our agent learning game variants in isolation. On this basis, we then introduce a relative scoring scheme which enables conclusions to be drawn at the benchmark level.

4.1 TRAINING VARIANT-EXPERTS

We aim to answer the following questions: (1) Does our agent achieve similar levels of performance when learning variants of the same game from scratch? (2) If not, what explains differences in performance across variants? In [Figure 2](#) we report means and standard deviations of top three variant-experts trained from scratch for 200 million environment steps. Our scores on default games are largely in line with those reported by other implementations of basic IQN agents, e.g. [Castro et al. \(2018\)](#). On remaining variants, we find that Rainbow-IQN performance varies widely for some game titles, less so on Freeway, even with independent hyper-parameter tuning. Overall, our agent’s scores are below variant performance ceilings, suggesting that some game variations may be harder to learn. While this may be at times due to the agent’s inductive biases—in particular, invisible objects—it is unlikely to fully explain observed variation, because such changes to basic games do not always have a detrimental impact on their own. For example, experts achieve significantly different levels of performance on identical variants of Space Invaders, except for visible vs invisible aliens (0.00 vs.0.08). However, experts have virtually equal performance with the same change to environment observations across Breakout variants, visible vs invisible walls of bricks (e.g. 0.00 vs. 0.12). Rather, we hypothesise that some variants are inherently harder to learn for the chosen agent, and thus a fixed training budget leads to performance differences across variations. This imposes a challenging bottleneck for transfer learning, which places high priority on limiting the resources expended for acquiring behaviours which maximise cumulative returns.

Statistical Analyses. We verify that variants have meaningful impact on learning by rejecting the null hypothesis that expert performance is the same across all combinations of the design factors which define game variants. We perform multi-factor Analysis of Variance (ANOVA) tests separately for each game title, and report results in [Table 1](#). Interaction effects are statistically significant in all cases, supporting the hypothesis that—although conceptually similar—design factors introduce significant changes to agent learning dynamics, even within the same game. In the post-hoc




Figure 2: Variant-Expert final raw score distributions, means and standard deviations across the top three experts in their respective hyper-parameter grids, reported for Space Invaders (top), Breakout (bottom left) and Freeway variants (bottom right). In all plots, hatching denotes the “difficulty” switch being activated, and colours indicate the “game mode”. See factorial design matrices for meanings of colour groups (Figure 1).

Table 1: Multi-Factor ANOVA (type 3) of Expert Scores vs. design factors.

Space Invaders	<i>df</i>	<i>F</i>	<i>PR(> F)</i>	Breakout	<i>df</i>	<i>F</i>	<i>PR(> F)</i>	Freeway	<i>df</i>	<i>F</i>	<i>PR(> F)</i>
Difficulty	1.0	115.90	$5.32e-16^a$	Difficulty	1	56.28	$1.26e-09^a$	Difficulty	1	1.39	$2.46e-01$
Invisible Invaders	1.0	1000.59	$8.67e-41^a$	Rules	2	1368.66	$4.71e-43^a$	Speeds	1	91.38	$6.75e-11^a$
Fast Bombs	1.0	314.09	$2.24e-26^a$	Extras	3	179.84	$4.22e-26^a$	Traffic	3	453.84	$2.73e-26^a$
Zigzagging Bombs	1.0	12.08	$9.22e-04^a$	<i>Interaction</i>	17	23.04	$1.91e-17^b$	<i>Interaction</i>	10	274.76	$4.42e-28^b$
Moving Shields	1.0	125.51	$9.93e-17^a$								
<i>Interaction</i>	26.0	1341.37	$5.33e-78^b$								

^astatistically significant after appropriate Bonferroni correction.

^bstatistically significant interaction effect at level: $\alpha = 0.05$.

analyses of differences between groups we found that all factors have a statistically significant effect on agent learning from scratch apart from the “difficulty” factor in Freeway. Further details can be found in Appendix C.

Discussion. Some differences in agent performance across variants were expected due to the qualitative impact of some factors on human play, e.g. changes in the dynamics of state transitions between “Breakout” and “Breakthru” variants, changes in observations such as visible vs. invisible invaders, or the effect of the difficulty switch. It is somewhat surprising that agents learn Freeway variants to similar performance, irrespective of difficulty level. Generally, we find that combinations of factors seem to induce substantially different learning dynamics, each having a significant effect on the final performance of the agent, at least within the standard training budget considered here. Hence, Rainbow-IQN demonstrates different levels of data efficiency across variants, despite their many similarities. A good amount of the variation can be explained by game design factors, which further emphasises the value of this benchmark for systematic and unbiased investigations of transfer with reinforcement learning algorithms.

Considering these empirical findings, as well as the possibility of different asymptotic performance ceilings for different variants, the remainder of the paper reports results relative to variant-expert performance, in order to meaningfully compare transfer learning across variations and curricula on a unified scale. We stress that variant-expert agents are generally not near-optimal. The most surprising results were the comparatively low agent score on “Timed Breakout”

variants (orange bars in the bottom left plot of Figure 2), which have identical transition dynamics to regular “Breakout”, but observations differ in that time in seconds is displayed instead of current score. Overall, these findings imply that scores normalised by variant-expert performance can be higher than 100%. Nevertheless, a relative scoring scheme accounts for the main effects of factors of variation on our agent’s data efficiency while learning particular game variants. This is essential for meaningful comparisons of transfer learning within and across game curricula.




Figure 3: Zero-shot transfer performance of default game experts (0.00) on all other variants, plotted as normalised scores (percentages) of the final performance of respective variant-experts, reported for Space Invaders (top), Breakout (bottom left) and Freeway variants (bottom right). Hatching denotes the “difficulty” switch being activated, and colours indicate the “game mode”. Colour groups are meaningful with respect to the factorial design matrices of particular game variants (Figure 1).

Table 2: Multi-Factor ANOVA (type 3) of default game expert zero-shot transfer vs. design factors.

Space Invaders	df	F	PR(> F)	Breakout	df	F	PR(> F)	Freeway	df	F	PR(> F)
Difficulty	1	319.42	1.43e - 26 ^a	Difficulty	1	182.22	5.83e - 18 ^a	Difficulty	1	15.87	3.66e - 04 ^a
Invisible Invaders	1	208.13	8.68e - 22 ^a	Rules	2	476.93	2.14e - 32 ^a	Speeds	1	211.61	1.17e - 15 ^a
Fast Bombs	1	250.95	7.92e - 24 ^a	Extras	3	177.63	5.53e - 26 ^a	Traffic	3	1524.22	1.36e - 34 ^a
Zigzagging Bombs	1	6.51	1.32e - 02	Interaction	17	31.38	2.87e - 20 ^b	Interaction	10	100.01	3.14e - 21 ^b
Moving Shields	1	116.63	4.67e - 16 ^a								
Interaction	26	117.61	1.57e - 44 ^b								

^astatistically significant after appropriate Bonferroni correction.

^bstatistically significant interaction effect at level: $\alpha = 0.05$.

4.2 ZERO-SHOT TRANSFER USING DEFAULT GAME EXPERTS

Game variants go beyond defaults in ways which are interesting to humans, providing additional challenges by design. Players will often become proficient in the default game, and then attempt to master others; indeed, several game manuals explicitly recommend this approach. Since deep neural networks are known to generalise to similar inputs, and since variants reuse the same visual elements (game sprites), we would like to answer the following questions: (1) How much transfer can be achieved purely through the generalisation power of deep neural networks trained on default games (0.00)? (2) Do factors of variation explain the success of zero-shot transfer (Taylor & Stone, 2009)?

Default game experts are evaluated, without further training, on all remaining variants of their respective games, and normalised scores are reported in Figure 3. Means and standard deviations of zero-shot transfer performance are plotted as percentages of the scores achieved by respective variant-experts, which were trained from scratch using 200 million interaction steps. Rainbow-IQN policies are derived using ϵ -greedy action selection according to value function approximations output by deep neural networks. In the zero-shot transfer case, errors of such predictions are unbounded, but depend on the differences in reward functions and transition dynamics between the source and transfer target MDPs, as well as on the specific function approximation method used. In the case of deep neural networks we have no formal guarantees for the quality of prediction on environment states never encountered in the source task, especially for those which do not exist at all in the source. All things considered, a surprising amount of positive zero-shot transfer is observed in all game domains, due to the sheer generalisation power of deep neural networks. In some cases, e.g. “Timed Breakout” variants, zero-shot transfer results in superior policies to what variant-experts have learned, although data and computational budgets were identical. However, these are exceptions. In the majority of cases, learning on default games results in lower return policies on other variants. That said, the observed zero-shot transfer is not negligible, and it does not appear uniform across game variations, which we investigate next.

Statistical Analyses. In order to answer the second question we check whether zero-shot performance is statistically indistinguishable across game variants, our null hypothesis of interest. We proceed with a multi-factor ANOVA of the dependent variable, here raw zero-shot transfer score, as a function of factors of variation in each game domain. Results reported in Table 2 indicate a statistically significant interactions in all games, meaning that factors play a role in the level of zero-shot transfer observed. In post-hoc analyses we must correct for multiple comparisons, and hence significance at level $\alpha = 0.05$ in Space Invaders is achieved for a corrected threshold of $\alpha_c = 0.0015$ or lower. We find statistically significant effects of all factors except the introduction of “Zigzagging Bombs” in Space Invaders.

Discussion. Zero-shot transfer results provide evidence that the benchmark offers ample challenges for transfer learning approaches. We find the highest levels of zero-shot policy transfer from default Breakout agents to variants, occasionally with higher performance than variant-experts. However, transferred policies are generally far from optimal, which naturally motivates investigations into fast adaptation through transfer learning. Overall, we observe substantial variance in the performance of fixed expert policies trained on defaults, then directly evaluated on all other variations. Although game variants within curricula are related, the statistically significant interaction effects between factors suggest that they are indeed distinct tasks with substantial and diverse “transfer gaps”, necessitating additional learning on top of default game policies. The most significant factors impacting zero-shot transfer were the “difficulty” switch in Space Invaders, the point accumulation “Rules” factor in Breakout, and “Traffic” thickness in Freeway, which we interpret as introducing the broadest differences. It is important to note that such modifications to default games do not easily fit within the more particular sets of assumptions often studied in transfer learning research. It is hard to argue that environment dynamics are precisely the same, or that state spaces are identical in distinct variants. Furthermore, even when observations change in subtle ways, e.g. the “difficulty” switch increases the size of the player’s cannon in Space Invaders, this change has a measurable impact on agent transfer performance, and it is not immediately clear how domain adaptation approaches would mitigate the issue. Hence, we go on to investigate how these diverse “transfer gaps” influence the performance of value-function finetuning, a more general transfer learning algorithm.

4.3 FINETUNING OF DEFAULT GAME EXPERTS

Zero-shot policy transfer can be substantial, and as such, constitutes an appropriate starting point for fine-tuning on other variants. However, it is not immediately obvious that current RL agents can capitalise on this zero-shot policy transfer to meaningfully speed up learning relative to variant-experts. In order to improve performance on new variants it is important that agent representations and predictions are finetuned on data from transfer target tasks. We report results of default game expert finetuning experiments in Figure 4, together with the natural baseline of zero-shot policy transfer from the same experts, and an ablation: learning from scratch for the same amount of environment interaction as used for finetuning. In a majority of cases we find that default expert finetuning outperforms baseline and ablation approaches. Hence, a simple and general algorithm such as Rainbow-IQN finetuning can leverage knowledge acquired in default games to improve performance beyond the zero-shot generalisation offered by its function approximator alone. Note that these gains are significant on the normalised scales induced by variant-expert performance, which account for the different learning dynamics caused by particular game variation, as shown in previous sections.

It is important not to overlook the heterogeneity of these finetuning results within and across curricula. Using a fixed, somewhat generous amount of data to finetune (10 million steps) does allow a substantial amount of performance to be acquired in some cases. Rainbow-IQN agents can also learn to a significant level of performance starting from scratch in this restricted training regime, which we would not have observed without the ablation experiment, e.g. on Freeway variants, but to some extent also on Breakout. Conversely, direct evaluation without training on the transfer target task can confer a competitive level of zero-shot performance, which we necessarily investigated in the previous section.




Figure 4: Finetuning performance starting from default game expert (0_00) on all other variants, plotted as normalised scores (percentages) of the final performance of respective variant-experts, reported for Space Invaders (top), Breakout (bottom left) and Freeway variants (bottom right). Please note that 100% indicates comparable performance with variant-experts, which were trained from scratch with $20\times$ more data than finetuning agents. Black bars show an ablation of finetuning training, starting from scratch rather than from default agent parameters. The delta to is due to transfer learning via agent network parameter initialisation. White bars indicate zero-shot transfer performance of default game agents. The delta to finetuning results represents the net benefit of learning in the transfer target variant. Hatching denotes the “difficulty” switch being on, and colours indicate the “game mode”, see Figure 1 for details.

Taken together, we interpret these heterogeneous finetuning results as strong motivation for research on improved general transfer learning algorithms, applicable across many transfer scenarios. But what can help transfer in this very general setting? Intuitively, learning several variants within curricula could offer a solution in the form of transfer using multiple sources, perhaps at the cost of additional complexity.

4.4 ZERO-SHOT TRANSFER SOURCE GAME SELECTION

In this section, we focus on knowledge in the form of policies learned by experts on other game variants, and we perform an all-to-all zero-shot transfer evaluation, reported as transfer matrices in Figure 5. We observe substantial off-diagonal structure in all games. The default game is not the best single source task in many cases, especially for variants with many factors introducing changes to the basic game. Naturally, we would like to know how well a superior source task selection strategy could perform on average. More precisely, we would like to empirically characterise the average upper bound on zero-shot transfer possible by appropriate selection of a single source policy for any game variant from all the remaining variant-expert policies. We plot the expectation for success of several transfer source task selection strategies in Figure 6, as well as the “best” case selection strategy, which provides an upper bound for what can be achieved by selecting a single variant-expert’s policy. For Freeway we observe strong transfer between policies learned with and without the “difficulty” switch being turned on. Matching the “best” strategy may be costly, e.g. in terms of expert evaluations, so we also report a “top3” selection strategy: zero-shot transfer performance of acting in every episode with one of the “top3” variant-experts at random. This strategy results in over 40% median normalised performance for all curricula. This level of transfer could be achieved at the expense of some target task interactions for ascertaining the quality of all other variant-expert policies on the current game variation. Interestingly, for all game titles, the “top3” strategy is superior to zero-shot transfer from the default games. This cannot be said for the average or “random” choice strategy, which is inferior to zero-shot transfer from default

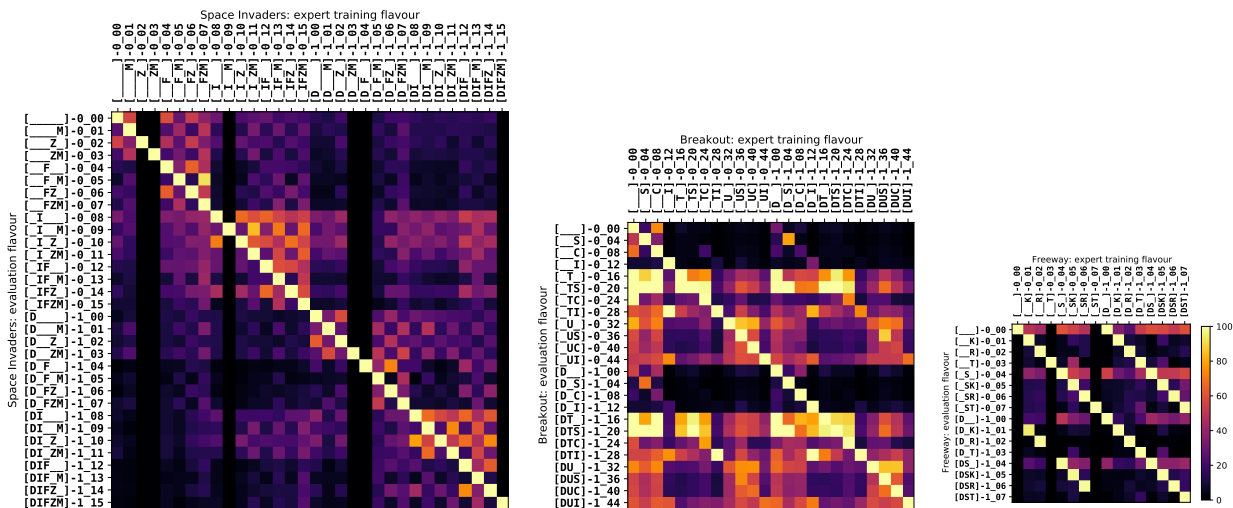


Figure 5: Transfer matrices depicting zero-shot evaluation of Space Invaders (left), Breakout (middle) and Freeway variant-experts (right), plotted as columns, and evaluated on all variants (rows). Values reported as normalised scores (percentages) of respective variant-expert scores.




Figure 6: Box-plots of zero-shot normalised score distributions for source task selection strategies on Space Invaders (left), Breakout (middle) and Freeway variants (right): “default” indicates using the default game expert (0.00) policy to act in all variants; the following selection strategies denote ways of choosing policies exclusively from *all other variants* in their respective domains: “random” means using a randomly chosen expert policy in every evaluation episode; “top3” and “best” strategies require access to some approximation of the transfer matrix to select the top three or best single source policy for each variant, from all other experts.

experts in Breakout. These observations motivate further research into automatic selection of variant-expert policies for transfer. The rich off-diagonal structure also recommends the use of Atari game variations as systematic curricula in a continual learning setting, if only for leveraging the potential for policy reuse highlighted by our evaluations.

5 CONCLUSIONS

Using Atari curricula, we have demonstrated that general transfer learning approaches, such as zero-shot policy generalisation and single expert model finetuning, lead to significant performance improvements in a large number of diverse transfer learning scenarios that are non-trivial for human players. Our analyses indicate that introducing modifications to basic games in a combinatorial fashion can substantially impact the empirical learning efficiency of a deep RL agent. Furthermore, zero-shot transfer performance from default games alone is also hindered by interactions between modifications. These results highlight the value of learning several variant-experts and reusing their knowledge to quickly find effective policies for new variations. Such strategies have the potential to efficiently match expert performance on transfer target tasks, but require the ability to both accumulate knowledge across game variants, and to appropriately transfer from several source tasks at once. Our study serves as a proof-of-concept that Atari curricula offer an invaluable benchmark of suitable complexity for systematically probing transfer without task engineering.

REFERENCES

- Andre Barreto, Will Dabney, Remi Munos, Jonathan J Hunt, Tom Schaul, Hado P van Hasselt, and David Silver. Successor features for transfer in reinforcement learning. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/file/350db081a661525235354dd3e19b8c05-Paper.pdf>.
- André Barreto, Diana Borsa, Shaobo Hou, Gheorghe Comanici, Eser Aygün, Philippe Hamel, Daniel Toyama, Shibl Mourad, David Silver, Doina Precup, et al. The option keyboard: Combining skills in reinforcement learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- 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*, 47:253–279, jun 2013.
- Marc G Bellemare, Will Dabney, and Rémi Munos. A distributional perspective on reinforcement learning. In *International Conference on Machine Learning*, pp. 449–458. PMLR, 2017.
- Richard Bellman. A markovian decision process. *Journal of mathematics and mechanics*, pp. 679–684, 1957.
- Yoshua Bengio. Deep learning of representations for unsupervised and transfer learning. In Isabelle Guyon, Gideon Dror, Vincent Lemaire, Graham Taylor, and Daniel Silver (eds.), *Proceedings of ICML Workshop on Unsupervised and Transfer Learning*, volume 27 of *Proceedings of Machine Learning Research*, pp. 17–36, Bellevue, Washington, USA, 02 Jul 2012. PMLR. URL <https://proceedings.mlr.press/v27/bengio12a.html>.
- Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemysław Debiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, et al. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*, 2019.
- Rich Caruana. Multitask learning. *Machine learning*, 28(1):41–75, 1997.
- 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.
- Karl Cobbe, Oleg Klimov, Chris Hesse, Taehoon Kim, and John Schulman. Quantifying generalization in reinforcement learning. In *International Conference on Machine Learning*, pp. 1282–1289. PMLR, 2019.
- Karl Cobbe, Chris Hesse, Jacob Hilton, and John Schulman. Leveraging procedural generation to benchmark reinforcement learning. In *International conference on machine learning*, pp. 2048–2056. PMLR, 2020.
- Will Dabney, Georg Ostrovski, David Silver, and Rémi Munos. Implicit quantile networks for distributional reinforcement learning. In *International conference on machine learning*, pp. 1096–1105. PMLR, 2018.
- Russell Davidson, James G MacKinnon, et al. *Estimation and inference in econometrics*, volume 63. Oxford New York, 1993.
- Jesse Farebrother, Marlos C Machado, and Michael Bowling. Generalization and regularization in dqn. *arXiv preprint arXiv:1810.00123*, 2018.
- Meire Fortunato, Mohammad Gheshlaghi Azar, Bilal Piot, Jacob Menick, Matteo Hessel, Ian Osband, Alex Graves, Volodymyr Mnih, Remi Munos, Demis Hassabis, et al. Noisy networks for exploration. In *International Conference on Learning Representations*, 2018.
- Ellen R Girden. *ANOVA: Repeated measures*. Number 84. Sage, 1992.
- Gene V Glass and Kenneth D Hopkins. *Statistical methods in education and psychology*. 1996.
- 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*, pp. 1861–1870. PMLR, 2018.
- Danijar Hafner, Timothy P Lillicrap, Mohammad Norouzi, and Jimmy Ba. Mastering atari with discrete world models. In *International Conference on Learning Representations*, 2020.

- 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 *Thirty-second AAAI conference on artificial intelligence*, 2018.
- Geoffrey E. Hinton, Simon Osindero, and Yee Whye Teh. A fast learning algorithm for deep belief nets. *Neural Computation*, 18:1527–1554, 2006.
- Dan Horgan, John Quan, David Budden, Gabriel Barth-Maron, Matteo Hessel, Hado van Hasselt, and David Silver. Distributed prioritized experience replay. In *International Conference on Learning Representations*, 2018.
- Timothy M. Hospedales, Antreas Antoniou, Paul Micaelli, and Amos J. Storkey. Meta-learning in neural networks: A survey. *CoRR*, abs/2004.05439, 2020. URL <https://arxiv.org/abs/2004.05439>.
- Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Z Leibo, David Silver, and Koray Kavukcuoglu. Reinforcement learning with unsupervised auxiliary tasks. *arXiv preprint arXiv:1611.05397*, 2016.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR (Poster)*, 2015.
- Sergey Levine and Vladlen Koltun. Guided policy search. In Sanjoy Dasgupta and David McAllester (eds.), *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pp. 1–9, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR. URL <https://proceedings.mlr.press/v28/levine13.html>.
- Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research*, 17(1):1334–1373, 2016.
- Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- 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.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidfjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.
- Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pp. 1928–1937. PMLR, 2016.
- OpenAI. Ai and compute, 2018. URL <https://openai.com/blog/ai-and-compute/>.
- Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2009.
- Emilio Parisotto, Lei Jimmy Ba, and Ruslan Salakhutdinov. Actor-mimic: Deep multitask and transfer reinforcement learning. In *ICLR (Poster)*, 2016.
- Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. Understanding the exploding gradient problem. *CoRR*, abs/1211.5063, 2012. URL <http://arxiv.org/abs/1211.5063>.
- Martin L Puterman. Markov decision processes: Discrete stochastic dynamic programming, 1994.
- Martin Riedmiller. Neural fitted q iteration—first experiences with a data efficient neural reinforcement learning method. In *European conference on machine learning*, pp. 317–328. Springer, 2005.
- Andrei A Rusu, Sergio Gomez Colmenarejo, aglar Gulehre, Guillaume Desjardins, James Kirkpatrick, Razvan Pascanu, Volodymyr Mnih, Koray Kavukcuoglu, and Raia Hadsell. Policy distillation. In *ICLR (Poster)*, 2016a.

- Andrei A. Rusu, Neil C. Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *CoRR*, abs/1606.04671, 2016b. URL <http://arxiv.org/abs/1606.04671>.
- Andrei A. Rusu, Matej Večerík, Thomas Rothörl, Nicolas Heess, Razvan Pascanu, and Raia Hadsell. Sim-to-real robot learning from pixels with progressive nets. In Sergey Levine, Vincent Vanhoucke, and Ken Goldberg (eds.), *Proceedings of the 1st Annual Conference on Robot Learning*, volume 78 of *Proceedings of Machine Learning Research*, pp. 262–270. PMLR, 13–15 Nov 2017. URL <https://proceedings.mlr.press/v78/rusu17a.html>.
- Tom Schaul, Daniel Horgan, Karol Gregor, and David Silver. Universal value function approximators. In Francis Bach and David Blei (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 1312–1320, Lille, France, 07–09 Jul 2015. PMLR. URL <https://proceedings.mlr.press/v37/schaul15.html>.
- Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay. In *ICLR (Poster)*, 2016.
- Tom Schaul, Hado van Hasselt, Joseph Modayil, Martha White, Adam White, Pierre-Luc Bacon, Jean Harb, Shihab Mourad, Marc G. Bellemare, and Doina Precup. The barbados 2018 list of open issues in continual learning. *CoRR*, abs/1811.07004, 2018. URL <http://arxiv.org/abs/1811.07004>.
- Simon Schmitt, Jonathan J Hudson, Augustin Zidek, Simon Osindero, Carl Doersch, Wojciech M Czarnecki, Joel Z Leibo, Heinrich Kuttler, Andrew Zisserman, Karen Simonyan, et al. Kickstarting deep reinforcement learning. *arXiv preprint arXiv:1803.03835*, 2018.
- Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, et al. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 588(7839):604–609, 2020.
- John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In Francis Bach and David Blei (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 1889–1897, Lille, France, 07–09 Jul 2015. PMLR. URL <https://proceedings.mlr.press/v37/schulman15.html>.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- David Silver, Satinder Singh, Doina Precup, and Richard S. Sutton. Reward is enough. *Artificial Intelligence*, 299:103535, 2021. ISSN 0004-3702. doi: <https://doi.org/10.1016/j.artint.2021.103535>. URL <https://www.sciencedirect.com/science/article/pii/S0004370221000862>.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- Matthew E Taylor and Peter Stone. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10(7), 2009.
- Sebastian Thrun and Lorien Pratt. *Learning to Learn: Introduction and Overview*, pp. 3–17. Springer US, Boston, MA, 1998. ISBN 978-1-4615-5529-2. doi: [10.1007/978-1-4615-5529-2_1](https://doi.org/10.1007/978-1-4615-5529-2_1). URL https://doi.org/10.1007/978-1-4615-5529-2_1.
- Tijmen Tieleman and Geoffrey Hinton. Rmsprop: Divide the gradient by a running average of its recent magnitude. coursera: Neural networks for machine learning. *COURSERA Neural Networks Mach. Learn.*, 2012.
- Joshua Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. *CoRR*, abs/1703.06907, 2017. URL <http://arxiv.org/abs/1703.06907>.

- Marin Toromanoff, Émilie Wirbel, and Fabien Moutarde. Is deep reinforcement learning really superhuman on atari? *CoRR*, abs/1908.04683, 2019. URL <http://arxiv.org/abs/1908.04683>.
- Hado Van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30, 2016.
- Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.
- Ziyu Wang, Victor Bapst, Nicolas Heess, Volodymyr Mnih, Remi Munos, Koray Kavukcuoglu, and Nando de Freitas. Sample efficient actor-critic with experience replay. 2016.
- Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine learning*, 8(3):279–292, 1992.
- Christopher John Cornish Hellaby Watkins. Learning from delayed rewards. 1989.
- Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on Robot Learning*, pp. 1094–1100. PMLR, 2020.
- Zhuangdi Zhu, Kaixiang Lin, and Jiayu Zhou. Transfer learning in deep reinforcement learning: A survey. *CoRR*, abs/2009.07888, 2020. URL <https://arxiv.org/abs/2009.07888>.

A APPENDIX: EXTENDED BACKGROUND

Q-learning. Watkins (1989) propose an RL algorithm which iteratively improves a parametric estimate $\hat{Q}_\theta = \hat{Q}_{\mathcal{M}}(\cdot, \cdot, \theta)$ of the optimal action-value function $Q_{\mathcal{M}}^*$ using the Bellman optimality operator (Bellman, 1957):

$$\hat{Q}_\theta(x, a) \leftarrow \mathbb{E}_{\mathcal{T}} [\mathcal{R}(x, a, x')] + \gamma \mathbb{E}_{\mathcal{T}} \left[\max_{a' \in \mathcal{A}} \hat{Q}_\theta(x', a') \right].$$

In model-free RL we do not assume explicit knowledge of the transition function \mathcal{T} , but the above optimisation can be implemented using samples procured by acting greedily with respect to the estimate \hat{Q}_θ with probability $1 - \epsilon$, and choosing actions uniformly at random otherwise. Watkins & Dayan (1992) show that, under certain condition, this procedure converges to $Q_{\mathcal{M}}^*$, and an optimal policy can be derived as $\pi_{\mathcal{M}}^*(x) = \arg \max_{a \in \mathcal{A}} \hat{Q}_\theta(x, a)$.

Deep Q-Learning. Mnih et al. (2015) use a convolutional neural network to parameterise \hat{Q}_θ , and train the resulting Deep Q-Network (DQN) by iterative minimisation of the squared temporal difference (TD) error:

Samples (x_t, a_t, r_t, x_{t+1}) are drawn randomly from a moving-window replay memory, where experience is collected online, and θ^- denotes a slow moving copy of online parameters θ . DQN agents are efficiently trained using RMSProp (Tieleman & Hinton, 2012), a variant of mini-batch stochastic gradient descent (SGD). This approach, with fixed hyper-parameters, was the first to master many classic Atari 2600 games from the Arcade Learning Environment (ALE) (Bellemare et al., 2013) up to a casual human level of performance, using screen pixels as inputs. Several improvements are brought together into the Rainbow agent by Hessel et al. (2018). Dabney et al. (2018) introduce Implicit Quantile Networks (IQN), a principled generalisation of DQN such that the full return distribution $Z_{\mathcal{M}}^{\pi}$ is modelled implicitly through its quantile function, leading to superior performance. Rainbow-IQN was shown to be state-of-the-art at the time for the standard (200 million steps) data regime by Toromanoff et al. (2019). This family of algorithms remains the most widely implemented of deep RL approaches.

Deep Reinforcement Learning. Online model-free RL with deep neural networks has been known to suffer from learning instability due to correlations between parameter updates. Stabilisation heuristics proved successful early on, such as off-line, full-batch training (Riedmiller, 2005; Schulman et al., 2015), using a replay memory to decorrelate parameter updates (Mnih et al., 2013; Schaul et al., 2016), and correcting over-estimation errors (Van Hasselt et al., 2016). Mnih et al. (2016) introduced A3C, an online deep RL algorithm which leveraged parallel data generation for stable learning, further improved with auxiliary losses (Jaderberg et al., 2016), and a replay memory (Wang et al., 2016) for better sample efficiency. Deep RL has also been successfully applied to continuous control domains (Lillicrap et al., 2015; Schulman et al., 2015; 2017; Haarnoja et al., 2018). Model-based RL algorithms have proven successful in full information environments which require complex strategy, such as the game of Go (Silver et al., 2016; 2017), several board-games (Schrittwieser et al., 2020) and domains with limited data, such as robotic control (Levine & Koltun, 2013; Levine et al., 2016). Recently, model-based RL approaches with learned models have proven competitive with model-free algorithms on domains with high-dimensional observations, e.g. DreamerV2 (Hafner et al., 2020), and MuZero (Schrittwieser et al., 2020). More complex training strategies have also been successful, such as combining deep RL with supervised learning from human games (Silver et al., 2016), and using self-play, or multi-agent training, leading to human Grandmaster level performance in Starcraft II (Vinyals et al., 2019), and winning against world champions at an e-sports game (Berner et al., 2019). It is important to note that such results are not yet achievable without access to super-computing infrastructure, e.g. thousands of GPUs or TPUs, and complex software stacks for distributed training, with estimated costs in single digit millions of dollars (OpenAI, 2018). Even when such resources are available, agent training can still take many months, so detailed ablations or hyper-parameter tuning become prohibitively expensive (Berner et al., 2019). Hence, improving data and computational efficiency through principled algorithmic innovations is essential for state-of-the-art deep RL.

B APPENDIX: AGENT TRAINING DETAILS

We tuned relevant hyper-parameters for each game variant in order to avoid sub-optimal performance due to poor settings. Please note that we are the first to report Rainbow-IQN results on Atari variants, and hence we had only the DQN experiments of Farebrother et al. (2018) as a guide, but no results were reported for the standard 200 million steps training regime, and the two algorithms have many differences which can impact learning dynamics. Farebrother et al. (2018) also used weight decay and dropout to reduce parameter norms and improve the generalisation of trained agents, but final performance was negatively impacted. We share the intuition that models which are closer to a random initialisation have an increased chance to generalise to related tasks, and may even be easier to adapt to new variants during finetuning.

Table 3: Variant-Expert Rainbow-IQN hyper-parameters. When multiple values are specified, the final hyper-parameter was chosen via grid search. Note that each configuration was trained using two different random seeds.

Name	Value(s)	Units/Comments/Attributions
Reward Clipping Bound	1	Possible values $\{-1, 0, 1\}$ (Mnih et al., 2015).
Replay Capacity	1 000 000	Transitions. (Mnih et al., 2015).
Replay Initial Size	500 000	Transitions. (Hessel et al., 2018).
Initial Policy	$\epsilon = 1$	Replay memory initialised using random policy.
Initial Policy ϵ Annealing Steps	100 000	Linearly annealed to $\epsilon = 0.01$ (Mnih et al., 2015).
Behaviour Policy (ϵ -greedy)	$\epsilon = 0.01$	The same during training and evaluation.
Action Repeats	4	Environment steps. (Mnih et al., 2015).
History Length	4	Agent steps. (Mnih et al., 2015).
Discount Factor	$\gamma = 0.99$	Mnih et al. (2015).
Multi-step Bootstrap	$n = 3$	Hessel et al. (2018).
Double Q-Learning	True	Van Hasselt et al. (2016).
Prioritised Replay	True	Schaul et al. (2016); Hessel et al. (2018).
Prioritisation Exponent	0.6	Horgan et al. (2018).
Prioritisation Type	proportional	Hessel et al. (2018).
Prioritisation Importance Sampling	0.4	Horgan et al. (2018).
Parallel Environments	1	Toromanoff et al. (2019).
Batch size	$\{32, 64\}$	Mnih et al. (2015)
Optimiser	Clipped-SGD	Re-scaling gradients with norms larger than C .
Gradient Clipping Norm	$C = 10$	Largest admissible gradient norm.
Agent Steps per Update	$\{4, 8\}$	Agent steps (Mnih et al., 2015).
Learning Rate	$\{10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}\}$	Hessel et al. (2018).
Target Network Update Period	$\{1000, 10\ 000\}$	Agent steps (Castro et al., 2018).
Number of τ Samples	$N = 64$	Dabney et al. (2018).
Number of τ' Samples	$N' = 64$	Dabney et al. (2018).
Number of $\tilde{\tau}$ Samples (Policy)	$K = 64$	Dabney et al. (2018).

We replaced RMSProp (Hinton et al., 2006), used by DQN (Mnih et al., 2015), and Adam (Kingma & Ba, 2015), used by Rainbow (Hessel et al., 2018), with Stochastic Gradient Descent (SGD) and gradient norm clipping (Pascanu et al., 2012). Please note that the aforementioned optimisers can provide superior performance, but tend to optimistically increase learning rates when parameter updates are correlated. We used hyper-parameter grid search to choose appropriate learning rates, and we purposefully also including smaller than typical values in the grid; we also allowed larger batch sizes and less frequent parameter updates besides standard values, see Table 3 for details. To further avoid parameter norms increasing due to correlated updates at the start of training, we sampled 500 000 transitions in the replay memory before starting learning, using ϵ decay from fully random behaviour to an ϵ -greedy policy with $\epsilon = 0.01$. The intuition behind these choices is that sampling data using a diverse set of behaviours policies will reduce the chances of correlated updates and learning divergence early on in training.

We used a slightly modified hyper-parameter grid for expert finetuning, primarily aimed at finding appropriate learning rates and the optimal levels of gradient clipping, as shown in Table 4. Expert finetuning is performed using a total of 10 million environment steps, which is $20\times$ less experience and parameter updates compared to expert training from scratch. Hence, we allowed for more lenient gradient clipping, as well as disabling it altogether.

Our choice to use hyper-parameter grid searches for both expert training and finetuning has greatly increased the computational costs of this study. We have somewhat offset these costs by reducing the numbers of filters in the convolutional trunk of the agent network, to the same sizes used by Mnih et al. (2013), but we kept the number of convolutional layers at 3 (Mnih et al., 2015; Dabney et al., 2018). Other network parameters were set according to Dabney et al. (2018), see architecture details in Table 5. These modifications did not significantly impact agent performance, which is largely comparable to results reported by Castro et al. (2018) on default games.

Table 4: Rainbow-IQN expert finetuning hyper-parameters. When multiple values are specified, the final hyper-parameter was chosen via grid search. Note that each configuration was trained using two different random seeds.

Name	Value(s)	Units/Comments/Attributions
Reward Clipping Bound	1	Possible values $\{-1, 0, 1\}$ (Mnih et al., 2015).
Replay Capacity	1 000 000	Transitions. (Mnih et al., 2015).
Replay Initial Size	500 000	Transitions. (Hessel et al., 2018).
Initial Policy	$\epsilon = 1$	Replay memory initialised using random policy.
Initial Policy ϵ Annealing	100 000	Linearly annealed to $\epsilon = 0.01$ (Mnih et al., 2015).
Behaviour Policy (ϵ -greedy)	$\epsilon = 0.01$	The same during training and evaluation.
Action Repeats	4	Environment steps. (Mnih et al., 2015).
History Length	4	Agent steps. (Mnih et al., 2015).
Discount Factor	$\gamma = 0.99$	Mnih et al. (2015).
Multi-step Bootstrap	$n = 3$	Hessel et al. (2018).
Double Q-Learning	True	Van Hasselt et al. (2016).
Prioritised Replay	True	Schaul et al. (2016); Hessel et al. (2018).
Prioritisation Exponent	0.6	Horgan et al. (2018).
Prioritisation Type	proportional	Hessel et al. (2018).
Prioritisation Importance Sampling	0.4	Horgan et al. (2018).
Parallel Environments	1	Toromanoff et al. (2019).
Batch size	32	Mnih et al. (2015)
Optimiser	Clipped-SGD	Re-scaling gradients with norms larger than C .
Gradient Clipping Norm	$C \in \{0, 10, 40\}$	Largest admissible gradient norm. Setting $C = 0$ disables gradient clipping.
Agent Steps per Update	4	Agent steps (Mnih et al., 2015).
Learning Rate	$\{10^{-7}, 10^{-6}, 5 \times 10^{-6}, 10^{-5}, 5 \times 10^{-5}, 10^{-4}\}$	Hessel et al. (2018).
Target Network Update Period	10 000	Agent steps (Castro et al., 2018).
Number of τ Samples	$N = 64$	Dabney et al. (2018).
Number of τ' Samples	$N' = 64$	Dabney et al. (2018).
Number of $\tilde{\tau}$ Samples (Policy)	$K = 64$	Dabney et al. (2018).

Table 5: Rainbow-IQN Network Architecture.

Name	Value(s)
Convolution Channels	16, 32, 32
Convolution Filter Sizes	$8 \times 8, 4 \times 4, 3 \times 3$
Convolution Filter Stride	$4 \times 4, 2 \times 2, 1 \times 1$
Convolution Activations	ReLU, ReLU, Identity
Embedding Hiddens	1568
Embedding Activations	ReLU
Latent Dimension	64
Combined Hiddens	256, 18
Combined Activations	ReLU, Identity

C APPENDIX: STATISTICAL ANALYSES

We used generalised linear models, fit with robust covariance estimates (‘HC3’), to perform multi-factor Analyses of Variance (ANOVA). After rejecting null hypotheses of the omnibus tests, we used Bonferroni corrections for multiple comparisons to perform post-hoc analyses of differences between group means. The desired significance level was

$\alpha = 0.05$, which we corrected to $\alpha_c = 0.0015$ for Space Invaders, $\alpha_c = 0.0020$, for Breakout, $\alpha_c = 0.0031$ for Freeway, since the numbers of groups are 32, 24 and 16 respectively.

Table 6: Multi-Factor ANOVA (type 3) of Expert Scores vs. design factors, extended tables.

<i>Space Invaders</i>	<i>sum_sq</i>	<i>df</i>	<i>F</i>	<i>PR(> F)</i>
Intercept	$8.18e + 07$	1	10862.43	$3.67e - 73$
Difficulty	$8.73e + 05$	1	115.90	$5.32e - 16$
Invisible Invaders	$7.54e + 06$	1	1000.59	$8.67e - 41$
Fast Bombs	$2.37e + 06$	1	314.09	$2.24e - 26$
Zigzagging Bombs	$9.10e + 04$	1	12.08	$9.22e - 04$
Moving Shields	$9.46e + 05$	1	125.51	$9.93e - 17$
Interaction	$2.63e + 08$	26	1341.37	$5.33e - 78$
Residual	$4.82e + 05$	64		

<i>Breakout</i>	<i>sum_sq</i>	<i>df</i>	<i>F</i>	<i>PR(> F)</i>
Intercept	$4.06e + 06$	1	7623.13	$1.49e - 54$
Difficulty	$2.99e + 04$	1	56.28	$1.26e - 09$
Rules	$1.46e + 06$	2	1368.66	$4.71e - 43$
Extras	$2.87e + 05$	3	179.84	$4.22e - 26$
Interaction	$2.08e + 05$	17	23.04	$1.91e - 17$
Residual	$2.55e + 04$	48		

<i>Freeway</i>	<i>sum_sq</i>	<i>df</i>	<i>F</i>	<i>PR(> F)</i>
Intercept	27174.26	1	134697.26	$1.44e - 59$
Difficulty	0.28	1	1.39	$2.46e - 01$
Traffic	274.67	3	453.84	$2.73e - 26$
Speeds	18.44	1	91.38	$6.75e - 11$
Interaction	554.30	10	274.76	$4.42e - 28$
Residual	6.46	32		




Figure 7: Empirical distribution plots of residuals from Multi-Factor ANOVA of Expert Scores vs. design factors.

C.1 ASSUMPTIONS:

1. The observations in each group are independent of the observations in every other group. We trained agents independently of each other in all cases, and we selected hyper-parameters for each group (i.e. game variation) independently of others.
2. Normality of measurement errors. For a give game, if we were to re-run the hyper-parameter grid search many times, choose the top 3 agents and average their performance, we assume that such measurements would be normally distributed. We have no reasons to believe that the distribution of the top 3 would be affected by divergence of online learning, one of the likely failure cases. We investigated empirical distribution of residuals from our analyses in Figure 7. We observe some deviation from normality, hence it is best to only draw conclusions based on the strongest significant effects; note that deviations from normality are relevant

Table 7: Multi-Factor ANOVA (type 3) of Default game expert zero-shot transfer vs. design factors, extended tables.

<i>Space Invaders</i>	<i>sum_sq</i>	<i>df</i>	<i>F</i>	<i>PR(> F)</i>
Intercept	$3.33e + 06$	1	638.21	$5.35e - 35$
Difficulty	$1.66e + 06$	1	319.42	$1.43e - 26$
Invisible Invaders	$1.08e + 06$	1	208.13	$8.68e - 22$
Fast Bombs	$1.31e + 06$	1	250.95	$7.92e - 24$
Zigzagging Bombs	$3.39e + 04$	1	6.51	$1.32e - 02$
Moving Shields	$6.08e + 05$	1	116.63	$4.67e - 16$
Interaction	$1.59e + 07$	26	117.61	$1.57e - 44$
Residual	$3.33e + 05$	64		

<i>Breakout</i>	<i>sum_sq</i>	<i>df</i>	<i>F</i>	<i>PR(> F)</i>
Intercept	$1.36e + 06$	1	3436.57	$2.51e - 46$
Difficulty	$7.23e + 04$	1	182.22	$5.83e - 18$
Rules	$3.78e + 05$	2	476.93	$2.14e - 32$
Extras	$2.11e + 05$	3	177.63	$5.53e - 26$
Interaction	$2.12e + 05$	17	31.38	$2.87e - 20$
Residual	$1.90e + 04$	48		

<i>Freeway</i>	<i>sum_sq</i>	<i>df</i>	<i>F</i>	<i>PR(> F)</i>
Intercept	1385.27	1	2775.16	$1.14e - 32$
Difficulty	7.92	1	15.87	$3.66e - 04$
Traffic	2282.53	3	1524.22	$1.36e - 34$
Speeds	105.63	1	211.61	$1.17e - 15$
Interaction	499.23	10	100.01	$3.14e - 21$
Residual	15.97	32		

when effect sizes are very small or when the experimental design is not balanced (Glass & Hopkins, 1996); in our case we use 3 samples per combination of factors (game variation), which is a balanced design.

3. Similar group variances of measurements. We used heteroskedasticity robust variance calculations of type ‘HC3’ (Davidson et al., 1993) to account for observed differences in variance, which we believe are primarily due to sample size.

D APPENDIX: COMPLETE RESULTS TABLES

In Table 8, Table 9 and Table 10 we provide raw and normalised scores for experts. In Table 11, Table 12, Table 13 we list raw and normalised scores of variant-expert zero-shot transfer.

Table 8: Space Invaders raw and variant-expert normalisation scores in various transfer learning settings and their ablations.

Space Invaders (Difficulty & Game Mod)	Raw Scores				Variant-Expert Normalised Scores		
	Variant- Expert (100%)	Finetuning from Scratch	Zero-shot Expert (0.00)	Finetuned Expert (0.00)	Finetuning from Scratch (%)	Zero-shot Expert (0.00) (%)	Finetuned Expert (0.00) (%)
0_00	2167.74	326.68	2167.78	33090.00	15.07	100.00	1526.48
0_01	1656.22	244.79	427.31	8949.71	14.78	25.80	540.37
0_02	2250.00	275.85	1243.42	2820.79	12.26	55.26	125.37
0_03	2000.00	235.40	412.60	2465.37	11.77	20.63	123.27
0_04	1783.01	157.97	260.19	1581.55	8.86	14.59	88.70
0_05	1170.38	121.72	89.57	634.09	10.40	7.65	54.18
0_06	1354.09	110.09	280.22	1017.78	8.13	20.69	75.16
0_07	1025.23	128.67	136.02	593.75	12.55	13.27	57.91
0_08	838.10	350.32	278.12	669.79	41.80	33.18	79.92
0_09	875.00	274.31	147.58	597.76	31.35	16.87	68.31
0_10	851.50	307.05	237.02	601.44	36.06	27.84	70.63
0_11	821.58	198.41	178.30	540.59	24.15	21.70	65.80
0_12	921.15	175.94	125.16	559.70	19.10	13.59	60.76
0_13	830.55	179.48	61.88	443.98	21.61	7.45	53.46
0_14	719.11	144.40	117.15	512.50	20.08	16.29	71.27
0_15	697.57	118.31	64.77	357.77	16.96	9.28	51.29
1_00	1938.98	315.28	105.91	1499.32	16.26	5.46	77.33
1_01	1207.91	226.72	131.86	623.19	18.77	10.92	51.59
1_02	1893.35	229.66	174.02	1244.93	12.13	9.19	65.75
1_03	1250.00	183.12	164.62	604.48	14.65	13.17	48.36
1_04	1250.00	145.50	17.18	576.14	11.64	1.37	46.09
1_05	929.55	68.88	40.73	385.40	7.41	4.38	41.46
1_06	993.20	88.99	20.10	524.81	8.96	2.02	52.84
1_07	713.80	135.76	58.13	342.39	19.02	8.14	47.97
1_08	857.34	352.02	65.88	648.87	41.06	7.68	75.68
1_09	886.21	186.99	46.95	530.19	21.10	5.30	59.83
1_10	845.13	304.50	64.59	591.04	36.03	7.64	69.93
1_11	829.79	283.46	87.92	485.00	34.16	10.60	58.45
1_12	729.67	82.16	24.88	476.72	11.26	3.41	65.33
1_13	653.80	141.09	24.23	203.37	21.58	3.71	31.11
1_14	644.76	77.95	24.39	413.87	12.09	3.78	64.19
1_15	600.00	60.90	15.48	204.83	10.15	2.58	34.14

Table 9: Breakout raw and variant-expert normalisation scores in various transfer learning settings and their ablations.

Breakout (Difficulty & Game Mode)	Raw Scores				Variant-Expert Normalised Scores		
	Variant- Expert (100%)	Finetuning from Scratch	Zero-shot Expert (0.00)	Finetuned Expert (0.00)	Finetuning from Scratch (%)	Zero-shot Expert (0.00) (%)	Finetuned Expert (0.00) (%)
0_00	264.60	6.69	264.60	342.21	2.53	100.00	129.33
0_01	420.10	5.25	212.41	412.43	1.25	50.56	98.17
0_02	298.60	6.51	201.41	325.49	2.18	67.45	109.01
0_03	247.00	6.05	22.02	69.99	2.45	8.92	28.33
0_04	71.90	82.43	116.85	111.95	114.65	162.51	155.70
0_05	55.50	70.97	106.79	101.32	127.87	192.42	182.56
0_06	255.40	72.79	253.50	320.40	28.50	99.26	125.45
0_07	82.40	75.70	50.19	90.10	91.87	60.91	109.35
0_08	386.30	202.85	316.54	385.42	52.51	81.94	99.77
0_09	670.10	251.56	387.64	635.95	37.54	57.85	94.90
0_10	650.20	170.03	357.95	591.00	26.15	55.05	90.90
0_11	385.60	203.56	197.61	356.12	52.79	51.25	92.35
1_00	155.80	14.41	86.12	198.12	9.25	55.28	127.16
1_01	398.00	6.65	50.89	364.01	1.67	12.79	91.46
1_02	335.60	9.46	59.29	223.65	2.82	17.67	66.64
1_03	166.30	14.62	15.19	40.97	8.79	9.14	24.64
1_04	55.50	75.97	62.92	81.70	136.89	113.36	147.20
1_05	51.40	69.62	63.44	80.50	135.44	123.43	156.61
1_06	188.10	92.15	110.40	434.00	48.99	58.69	230.73
1_07	63.20	74.81	35.28	67.71	118.37	55.82	107.13
1_08	353.20	133.44	260.52	362.86	37.78	73.76	102.74
1_09	542.40	217.99	299.91	450.62	40.19	55.29	83.08
1_10	516.90	182.10	293.05	455.30	35.23	56.69	88.08
1_11	361.90	126.34	219.33	343.84	34.91	60.61	95.01

Table 10: Freeway raw and variant-expert normalisation scores in various transfer learning settings and their ablations.

Freeway (Difficulty & Game Mode)	Raw Scores				Variant-Expert Normalised Scores		
	Variant- Expert (100%)	Finetuning from Scratch	Zero-shot Expert (0.00)	Finetuned Expert (0.00)	Finetuning from Scratch (%)	Zero-shot Expert (0.00) (%)	Finetuned Expert (0.00) (%)
0_00	33.10	23.31	33.14	33.43	70.43	100.12	100.99
0_01	25.83	5.94	0.37	23.55	22.99	1.43	91.18
0_02	27.13	4.98	0.10	11.14	18.35	0.38	41.07
0_03	32.97	11.06	0.24	24.90	33.55	0.73	75.52
0_04	33.30	21.86	18.30	24.55	65.64	54.94	73.73
0_05	25.65	7.96	1.74	9.94	31.03	6.80	38.74
0_06	25.66	7.98	1.53	10.10	31.10	5.98	39.38
0_07	28.74	13.23	1.19	16.49	46.03	4.13	57.38
1_00	33.08	16.94	33.06	33.31	51.20	99.93	100.68
1_01	26.87	5.31	0.23	23.86	19.75	0.86	88.79
1_02	27.34	0.74	0.00	0.00	2.69	0.01	0.00
1_03	32.88	8.16	0.01	23.86	24.83	0.02	72.56
1_04	33.39	17.98	13.16	23.37	53.85	39.41	69.98
1_05	25.85	4.92	1.01	7.94	19.03	3.90	30.71
1_06	25.66	4.67	0.74	7.06	18.21	2.89	27.52
1_07	28.81	7.67	0.42	14.24	26.63	1.48	49.44

Table 13: Freeway zero-shot transfer matrix between variant-experts trained from scratch. Top: raw scores, Bottom: Variant-Expert normalised scores (percentages).

	0.00	0.01	0.02	0.03	0.04	0.05	0.06	0.07	1.00	1.01	1.02	1.03	1.04	1.05	1.06	1.07
0.00	33.10	16.24	14.70	n/a	20.03	18.25	15.86	n/a	33.14	11.71	8.44	17.48	19.60	19.06	17.14	19.80
0.01	n/a	25.83	1.95	n/a	1.25	4.49	3.23	n/a	0.37	26.68	2.63	1.89	0.56	4.86	1.53	2.15
0.02	n/a	2.66	27.13	n/a	1.32	1.47	3.41	n/a	0.10	1.87	27.41	2.65	0.71	1.50	2.31	3.94
0.03	n/a	2.99	0.91	32.97	2.28	14.01	0.73	n/a	0.24	3.34	4.36	32.97	3.25	7.72	0.07	1.43
0.04	n/a	13.07	11.38	n/a	33.30	17.76	15.65	n/a	18.30	10.37	8.20	15.01	33.47	16.98	17.25	18.92
0.05	n/a	3.31	3.60	n/a	3.03	25.65	5.07	n/a	1.75	1.95	2.53	4.22	2.32	25.82	3.95	8.26
0.06	n/a	3.52	4.27	n/a	3.51	5.82	25.66	n/a	1.54	1.95	3.50	3.13	2.40	5.47	25.34	7.13
0.07	n/a	3.90	4.41	n/a	1.15	4.53	0.61	28.74	1.19	2.55	4.01	8.34	1.05	2.96	0.58	28.74
1.00	n/a	6.90	6.54	n/a	16.35	11.84	9.87	n/a	33.08	5.63	2.87	7.51	14.04	12.37	10.61	13.04
1.01	n/a	25.50	0.79	n/a	1.21	2.73	2.44	n/a	0.23	26.87	2.13	0.32	0.47	3.18	0.78	0.67
1.02	n/a	0.26	27.00	n/a	0.06	0.07	0.59	n/a	n/a	0.75	27.34	0.20	0.31	0.10	0.48	0.13
1.03	n/a	1.98	0.78	n/a	0.10	6.19	0.86	n/a	0.01	1.01	1.28	32.88	0.23	7.99	0.04	0.12
1.04	n/a	8.46	6.47	n/a	32.97	13.40	11.94	n/a	13.16	5.86	5.28	10.39	33.39	12.73	12.56	13.73
1.05	n/a	1.83	1.76	n/a	1.87	25.27	4.01	n/a	1.01	0.90	1.04	2.13	1.60	25.85	3.37	4.65
1.06	n/a	2.04	2.62	n/a	2.45	2.94	25.39	n/a	0.74	1.01	2.42	1.45	1.70	3.41	25.66	4.44
1.07	n/a	1.73	1.84	n/a	0.75	3.25	0.35	n/a	0.42	1.18	2.17	4.87	0.45	2.14	0.37	28.81
	0.00	0.01	0.02	0.03	0.04	0.05	0.06	0.07	1.00	1.01	1.02	1.03	1.04	1.05	1.06	1.07
0.00	100.00	49.06	44.42	n/a	60.52	55.14	47.91	n/a	100.13	35.38	25.49	52.80	59.22	57.57	51.78	59.81
0.01	n/a	100.12	7.54	n/a	4.86	17.39	12.50	n/a	1.43	103.40	10.18	7.34	2.16	18.82	5.93	8.34
0.02	n/a	9.82	100.10	n/a	4.89	5.41	12.57	n/a	0.38	6.92	101.14	9.76	2.63	5.55	8.51	14.56
0.03	n/a	9.05	2.77	99.91	6.90	42.44	2.20	n/a	0.73	10.13	13.22	99.92	9.86	23.39	0.21	4.33
0.04	n/a	39.25	34.19	n/a	100.00	53.33	47.01	n/a	54.94	31.14	24.63	45.08	100.51	51.01	51.80	56.83
0.05	n/a	12.89	14.01	n/a	11.79	99.81	19.72	n/a	6.80	7.59	9.84	16.43	9.02	100.48	15.37	32.13
0.06	n/a	13.68	16.60	n/a	13.64	22.64	99.84	n/a	5.97	7.58	13.60	12.16	9.32	21.29	98.59	27.76
0.07	n/a	13.60	15.37	n/a	4.00	15.80	2.12	100.14	4.13	8.89	13.99	29.07	3.65	10.31	2.01	100.14
1.00	n/a	20.84	19.77	n/a	49.40	35.77	29.83	n/a	99.93	17.01	8.67	22.70	42.40	37.37	32.05	39.41
1.01	n/a	94.79	2.94	n/a	4.51	10.15	9.06	n/a	0.86	99.90	7.91	1.21	1.73	11.82	2.89	2.51
1.02	n/a	0.95	98.90	n/a	0.23	0.26	2.15	n/a	0.01	2.73	100.16	0.72	1.15	0.38	1.77	0.47
1.03	n/a	6.01	2.36	n/a	0.32	18.80	2.63	n/a	0.02	3.07	3.88	99.92	0.69	24.29	0.12	0.35
1.04	n/a	25.34	19.36	n/a	98.70	40.12	35.75	n/a	39.41	17.54	15.82	31.11	99.96	38.12	37.62	41.11
1.05	n/a	7.11	6.82	n/a	7.25	97.96	15.55	n/a	3.91	3.49	4.04	8.25	6.21	100.19	13.05	18.01
1.06	n/a	7.93	10.20	n/a	9.53	11.45	98.81	n/a	2.89	3.94	9.40	5.65	6.63	13.27	99.85	17.29
1.07	n/a	5.99	6.40	n/a	2.62	11.28	1.23	n/a	1.48	4.08	7.54	16.92	1.57	7.42	1.30	100.03