



ADVANCED

ANIMAL TRAINING

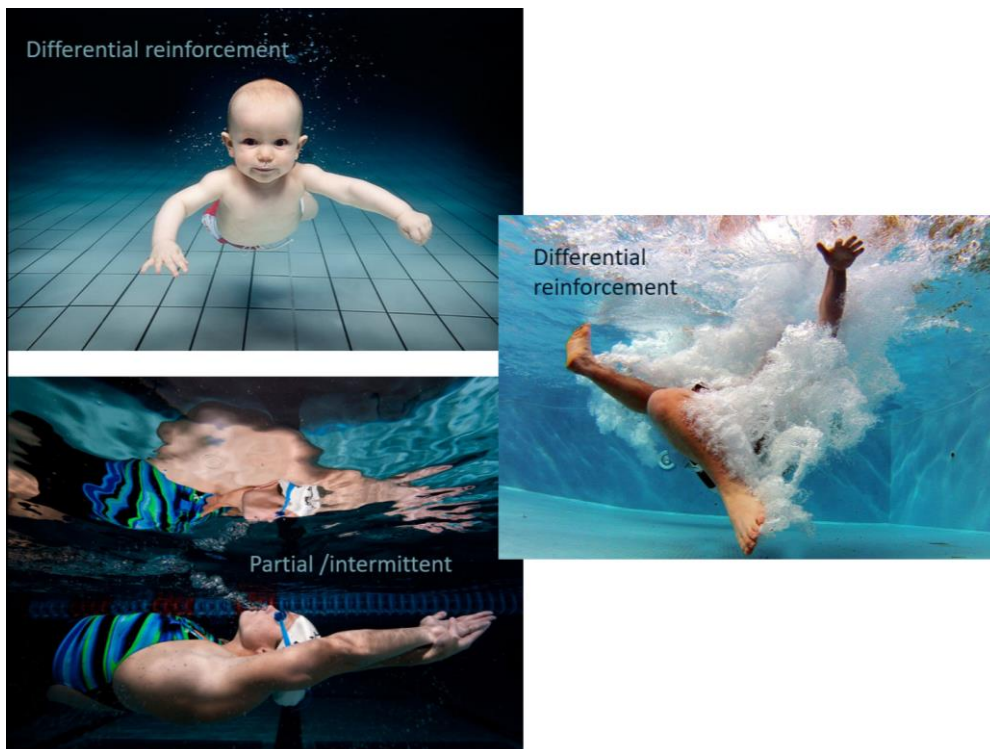
7. Schedules of Reinforcement.

Please don't distribute this material without my consent.

Chapter 1.

Reinforcement schedules are all about reinforcing some responses and not others.


They come in two families: differential reinforcement (typically used when teaching animals a new behavior or to drop an unwanted behavior), and partial or intermittent reinforcement (most often used to maintain already learned behavior).



Chapter 2.

Continuous reinforcement is avoided by some trainers, claiming that responses maintained by continuous reinforcement get too sensitive to extinction.

I disagree with that: resistance to extinction is built in as the animal learns the behavior – some responses are reinforced, others not. During shaping, only the best responses are selected for reinforcement.



Chapter 3.

Intermittent schedules may be viewed from two dimensions, related to how predictable the schedule is (fixed or variable) and the unit of measure (most typically response ratios or time intervals). These can be illustrated in a two-by-two matrix, producing four outcomes.

In fixed ratio schedules, the animal needs to perform a fixed number of responses before reinforcement occurs. In variable ratio schedules, that number is random.

In fixed interval schedules, the first behavior occurring after a specific time interval is reinforced. In variable intervals, that time interval is random.

Animals learn when to expect reinforcement depending on the schedule they're on, and behavior changes predictably.

	FIXED	VARIABLE
RATIO	FIXED RATIO SCHEDULE	VARIABLE RATIO SCHEDULE
INTERVAL	FIXED INTERVAL SCHEDULE	VARIABLE INTERVAL SCHEDULE


Chapter 4.

Variable ratio schedules (VR) leads to a high and even rate of responding.

Fixed ratio schedules (FR) lead to a rather high rate of responding, but with post-reinforcement pauses after each reinforcer is received. These pauses are imperceptible in small ratios and may be very long in large ratios.

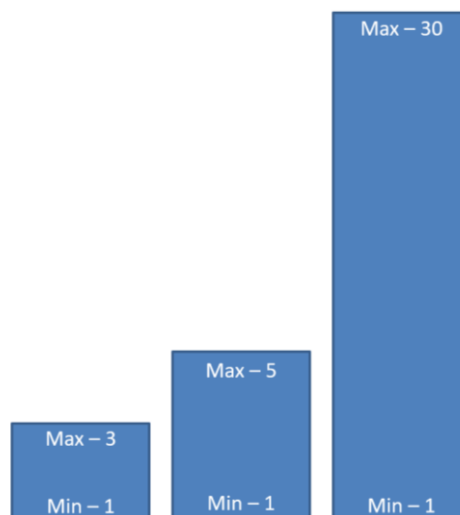
Variable interval schedules (V) lead to a low and even rate of responding.

Fixed interval schedules (FI) develop a typical scalloped pattern, when the animal is trying to estimate when the next opportunity for reinforcement arises, and increases response rates towards the end of each interval.



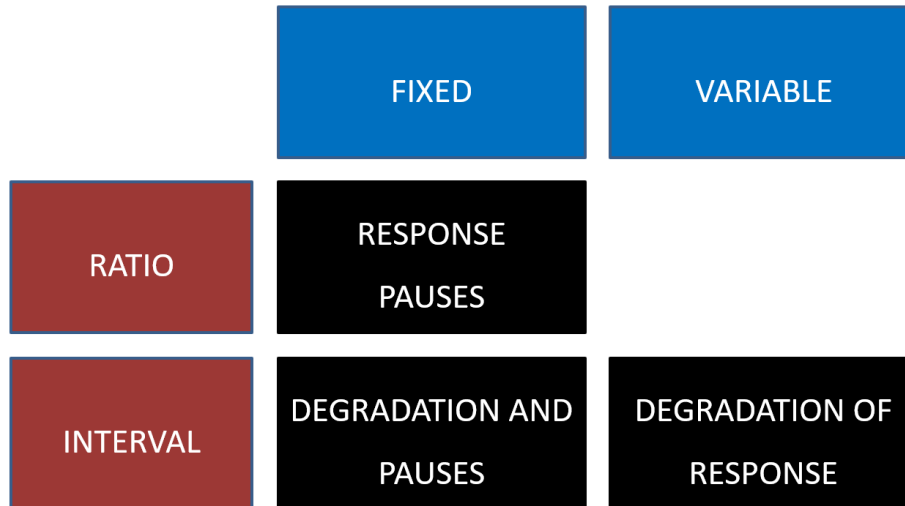
Chapter 5.

Implementing variable ratio schedules is tricky since if you stretch the ratio too soon you risk ratio strain: that the animal becomes frustrated or gives up. Typically, it's suggested to use low mean ratios, maintain 1 as the lowest possible ratio, and gradually increase the amplitude while staying at a low mean before increasing the mean.




Chapter 6.

Different intermittent schedules lead to different types of problems. One of the biggest mistakes people make is that they use them too early – before the behavior is fluent. Also, schedules may increase distractibility, so ask yourself if you really need an intermittent schedule before you teach it.



Chapter 7.

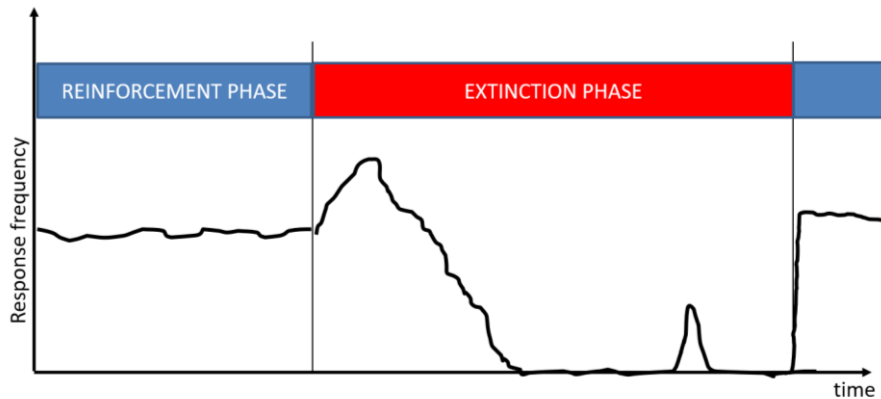
Using reinforcement variety may be a way to tap into the benefits of variable ratio schedules without the downsides. Just don't assume that you know what's reinforcing to a particular animal.



Chapter 8.

During extinction, behavior may go through a brief period when it's shown more frequently (called an extinction burst). There might also be resurgence of behavior: the sudden reinstatement of the behavior that's undergoing extinction.

It typically only takes a single reinforcer to bring the response back to the full intensity that it had before extinction occurred – many extinction schedules may turn into intermittent schedules, especially if you don't have full control of the reinforcer. That makes the response resistant to extinction.



Chapter 9.

There are three types of differential reinforcement schedules:

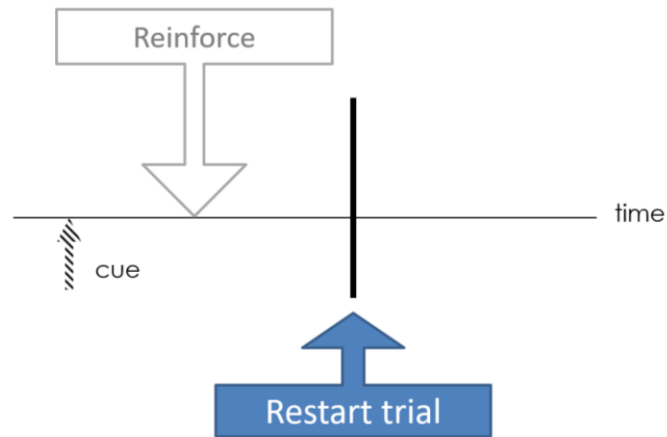
Duration schedules, where animals learn to keep performing a specific response until the schedule requirement is met.

Response type schedules, where only some types of responses are reinforced (DRE - excellent; DRA - alternative; DRO - other; DRI - incompatible), and response rate schedules, where only some rates of responses are reinforced (DRL - low, DRP - paced, DRH - high, DRO/Z - zero).



Chapter 10.

Limited holds are a procedure that trainers use to reduce the latency of responding. If the animal responds before the mean latency is reached, that response is reinforced. If the animal hasn't responded before the mean latency is reached, the trial is restarted. Then the mean latency is gradually shaped shorter.



Chapter 11.

Behavioural momentum is the animal trainer's analogy to momentum in physics. By using this analogy, we can get better compliance and build resistance to extinction – without resorting to intermittent schedules.

reinforcement = change mass x speed

Chapter 12.

Animals (and humans) will do more of the behaviour that leads to the most favoured outcome, but they will keep offering the other, less well reinforced behaviour too, at least sometimes.

Matching is affected by reinforcer quality, rate and delay of reinforcement - and response effort.

This means that if you want the animal to perform the behaviour leading to the least favoured outcome, you must increase the value of that outcome somehow.

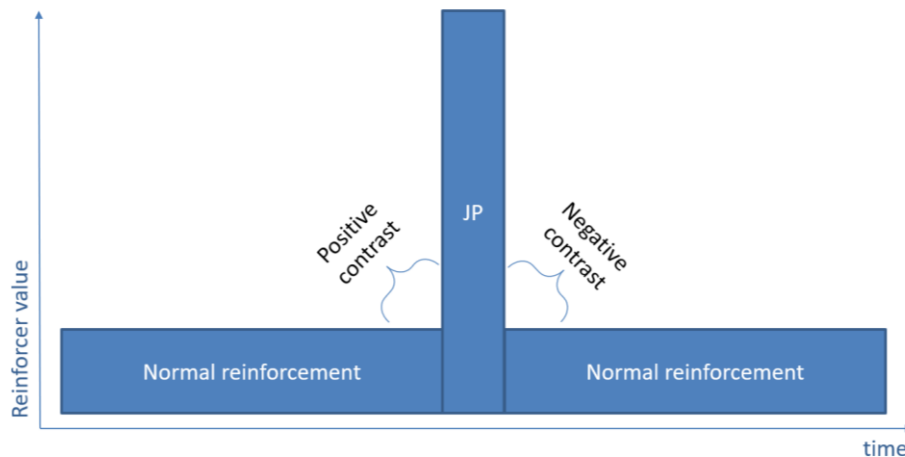


Chapter 13.

Jackpots are inconsistently used by trainers and their benefit is debated. Jackpots are unexpected, high-value reinforcers, sparingly used to pinpoint breakthroughs in training.

People aim for the positive contrast to the preceding behaviour, but forget about the negative contrast to the upcoming behaviour – which may be aversive.

Don't confuse jackpots with predictably paying more for difficult behaviour!



Chapter 14.

The Premack Principle suggests that animals are willing to do boring things (low-probability behaviour) in order to get the opportunity of doing fun things (high-probability behaviour).

This is used by contemporary animal trainers in two ways: to increase desired low-probability behaviour, and to control undesired high-probability behaviour.

Premack reinforcers may be difficult to use, since the animal may become aroused when engaged in the high-probability behaviour. It may be several minutes before the animal is calm enough to engage in the low-probability response again.

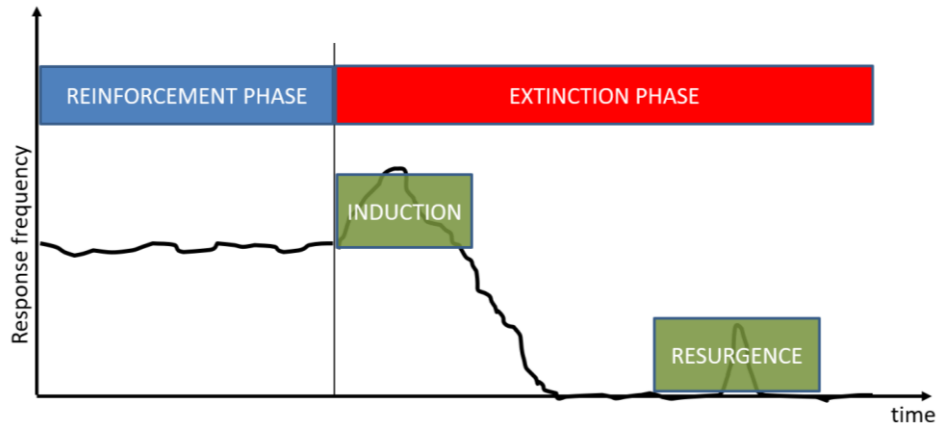
Premack principle – two uses

- Reinforce desired (low-probability) behaviour
- Manage undesired (high-probability) behaviour

Chapter 15.

Behaviour can be directly or indirectly reinforced. Indirect reinforcement is termed induction and occurs in three situations: for behaviour preceding directly reinforced behaviour, in between two directly reinforced behaviour, or the increased variability seen during an extinction burst.

Resurgence is the return of previously reinforced behaviour, which occurs during regression and extinction trials.



Chapter 16.

Perhaps one reason why there's confusion about intermittent schedules is because of the plethora of different – and contradictory – ways to depict them graphically...

