FACULTY OF SCIENCE
Dr Douglas Elliffe
Deputy Dean
Professor of Psychology

THE UNIVERSITY OF AUCKLAND
FACULTY OF SCIENCE

Office of the Dean of Science
Science Centre, 23 Symonds St
Auckland, New Zealand
Telephone 64 9 9235262
Facsimile 64 9 367 7111
Email: d.elliffe@auckland.ac.nz
www.science.auckland.ac.nz

The University of Auckland
Private Bag 92019
Auckland, 1142, New Zealand

26 November 2020

Commentary on China, L, Mills, D.S., & Cooper, J.J. (2020). Efficacy of dog training with and without electronic collars vs. a focus on positive reinforcement. *Frontiers in Veterinary Science*, 7:508. https://doi.org/10.3389/fvets.2020.00508

China, Mills and Cooper (2020)[1] compared a group of dogs trained using electronic shock collars[2] with two control groups trained without e-collars. The trained target behaviours were recall to the trainer in response to the command "Come" and sitting in response to the command "Sit". In terms of different measures of performance, they reported either no differences between groups or better performance for the control group trained without e-collars and with an explicit emphasis on positive reinforcement. China et al. interpreted these findings as evidence that e-collars should not be used in dog training.

However, the reported research suffers from (1) sufficient external and internal confounding variables to weaken severely any comparison between groups; (2) sometimes inappropriate data analyses; (3) a very strange choice of target behaviours to train, perhaps betraying an oversimplified understanding of the principles of behavioural psychology; (4) incautious conclusions that go well beyond the reported data, both in terms of the specific application of e-collars in the study, but more seriously in generalization to other possible uses of e-collars. Furthermore, the anticipated effect of each of these flaws is consistently in the direction unfavourable to a finding that e-collar use was effective.

My conclusion is that China et al., at least if interpreted much more cautiously than do its authors, makes a flawed but not entirely valueless contribution to the scientific literature.

[1] Henceforward 'China et al.'
[2] Henceforward 'e-collars'

However, as evidence to inform evidence-based policy making, it is very seriously flawed and should not be relied on. A move to ban e-collars based on China et al. would not be an evidence-based decision.

Confounding variables:

To provide useful information, an experimental design must take rigorous steps to rule out all possible explanations other than the independent variable under investigation. If this is not done, then any effect cannot be unequivocally attributed to the independent variable (here, e-collar use and/or an emphasis on positive reinforcement). These potential alternative explanations are called *confounding variables* or simply *confounds*. They come in two flavours – internal confounds concern pre-existing differences between the research participants, or between groups of participants subject to different treatments (different applications of the independent variable). External confounds concern the environment within which the data are collected. The aim is that all of the conditions of data collection are identical except for the different treatments. China et al.'s research includes obvious examples of both kinds of confound, with the potential to have had serious impact on their results.

Internal confounds:

China et al. chose a between-groups experimental design, in which different dogs were allocated to the three treatment groups. The immediate disadvantage of such a design is that there may be preexisting differences between groups (internal confounds), and the classic way to minimize this problem is by random allocation of participants to groups. This was not done – dogs were randomly allocated to the e-collar group[3] and one control group[4], but were recruited separately at a later date for the second control group[5] whose training emphasized positive reinforcement. The authors claimed that this conferred the advantage of being able to match the behavioural problems for which the dogs had been referred to training. They were unsuccessful in this – dogs with poor recall (i.e., poor response to a command to return to the owner) but no chasing (e.g., of sheep) issues were three times as common in Group C2 as in Group E, and dogs with a history of aggression against other dogs were twice as common in Group E as in Group C2. That is, the standard design technique (randomization)

[3] Henceforward 'Group E'
[4] Henceforward 'Group C1'
[5] Henceforward 'Group C2'

for between-groups comparisons was not employed, and in fact pre-existing differences appear to have emerged. Those differences were in the direction unfavourable to the e-collar – Group E appears likely to have been a more difficult group to train than Group C2 in particular.

China et al. disregard this as a factor on the grounds that there were no significant differences between groups in reasons for referral. This is a misunderstanding of the nature of statistical inference. A significant difference means a difference whose probability of occurring solely due to sampling error from a single population is less than a conventional value, usually .05. This is a deliberately conservative criterion to guard against the report of false positives – apparent differences between groups that are in fact simply due to chance – and has the effect that any difference must be quite large, relative to within-group variability, to entitle us to claim an effect. The converse is not true – a failure to find statistical significance only means that any difference is not large enough to justify the strong claim that there is an effect. 'Not significantly different' does not mean 'significantly the same', yet China et al. have behaved as if it does. Failure to find a statistically significant pre-existing difference between groups does not entitle one to make the strong claim that there was no pre-existing difference.

External confounds:

The usual advantage of a between-groups research design is that it allows external confounds to be minimized, and ideally eliminated. This is because there is a single measure of behaviour for each participant (or, as here, a single average of multiple measures), so that all of the conditions under which the data are collected can be kept constant. This was not done by China et al. – Groups E and C1 were trained on a farm in Scotland in early winter, sometimes in weather conditions sufficiently severe that training had to be moved indoors (further weakening standardization of conditions). But Group C2 was trained at a different location (a dedicated training centre) in a different part of the UK (Lincoln) and in a different and more clement season (spring). All these differences appeal as potential alternative explanations for greater training success for Group C2.

China et al. do note this concern but address it in a surprising, unconvincing and incomplete manner. They point to preliminary data analysis that showed no change across weeks of training within groups. The issue is, however, that there were differences in conditions

between groups, and therefore correlated with the independent variable (different treatments), and therefore potential alternative explanations of different behaviour between groups.

As with the internal confounds noted above, these are relatively minor weaknesses that, in a dispassionate scientific report, can be acknowledged as reasons to be cautious about implications of the results. To some extent, China et al. do so, although in my view not enough. They are not necessarily reasons not to publish the research, although a more prestigious journal than the *Frontiers* series would probably have required higher scientific standards. But they do suggest great caution in basing policy decisions on these results – the research is too flawed to do so safely.

There is, however, a much more serious external confound than these. Group E's training involved both the use of the e-collar and positive reinforcement in the form of food reward, and Group C2 received only positive reinforcement. But C2, according to the authors, received about 5 times as much positive reinforcement as E. Given that it is very well established that positive reinforcement is a useful tool for increasing a behaviour, this means that C2 was advantaged because it received more positive reinforcement. That is, E differed from C2 not only in use of the e-collar but also in another respect that would have affected behaviour. This means that the efficacy of the e-collar simply can't be assessed, because the comparison with C2 is confounded.

It is likely that China et al. don't recognize this as a problem, because they apparently set out to prove an advantage of an 'emphasis on reward-based training' (p. 3) and their curious focus is on comparing 'industry best-practice' training regimens. However, the question of e-collar efficacy is fundamentally unanswerable from this research, because the frequency of positive reinforcement is not equated between groups. It is perfectly plausible that if Group E had received the same frequency of positive reinforcement as Group C2, together with e-collar use as well, its performance may have been superior. It is only by holding all other aspects of the training constant that the efficacy of the e-collar can be assessed, and China et al. did not do that.

Noting the caveats around other internal and external confounds above, China et al.'s results provide some, fairly ambiguous, evidence that positive reinforcement alone is at least as effective in training two specific behaviours as is a combination of e-collar use and less

frequent positive reinforcement. What it does not do is dismiss the possibility that the e-collar would confer an advantage if all other aspects of training are held constant.

A final weakness in how training was conducted relates to the use of a long lead most of the time. There are at least three problems here. Firstly, time spent on the lead was entirely at the trainers' discretion, again weakening standardization of conditions between groups and introducing an external confound. Secondly, it is very likely that dogs can discriminate whether they are on a lead or not, and their behaviour is likely to differ depending on that. In particular, running away and failing recall is much more likely when they know that there is no lead preventing them from doing so. But most importantly, the dogs were within 1 m of the trainer about 70% of the time, and further than 5 m away less than 5% of the time. It is precisely when the dog is some distance away and not attached to a lead that a remotely controllable e-collar may be useful – again, the conditions under which the research was conducted are stacked against the e-collar.

Inappropriate data analysis:

Evaluating the efficacy of different training techniques requires examining how behaviour changed during training, not simply averaging group performance across all training stages. That is, we want to know whether the dogs are learning as a result of their training. From this perspective, there appeared to be little change in behaviour across training (Figs 4 and 5), suggesting that none of the training regimes was very effective, including Group C2. However, there is some suggestion (Fig 4) that latency to return decreased (improved recall) in Groups E and C1.

China et al. note this but report that the apparent improvement was not statistically significant. It is difficult to be certain, but this conclusion appears to have been reached from a between-groups ANOVA, effectively treating training day as a between-groups fixed factor. This is inappropriate – the same dogs' behaviour was measured on multiple occasions (days), so this aspect of the results is amenable to analysis using a more powerful repeated-measures approach. That is, it would have been possible to test for a trend of decreasing latency across days in the behaviour of each individual dog. This has the advantage of making between-dogs variability irrelevant to the analysis – each individual is compared with itself. Therefore, the analysis used is unsuitable for determining whether any of the training

procedures was effective in changing behaviour. Surely, producing change in behaviour is the point of any training procedure.

Choice of target behaviours:

To this point, I have concentrated on showing that China et al. have not succeeded in their aim of empirically comparing training including e-collars with training that only used positive reinforcement as ways of increasing obedience to the commands 'Come' and 'Sit'. They have not succeeded because of flaws in their research design and, to a lesser extent, their data analysis. Even if they had done this successfully, however, I now argue that their aim was the wrong one, and that their choice of behaviours to train was similarly difficult to justify.

A brief digression into how behaviour is controlled by its consequences is necessary at this point. The fundamental principle of behavioural psychology is the Law of Effect. This states that when behaviour is followed by a desirable consequence (*reinforcement*) it will be strengthened and increase in future likelihood. If it is followed by an undesirable consequence (*punishment*) it will be weakened and decrease in likelihood. Both reinforcement and punishment can be arranged using either a positive or a negative contingency[6]. In *positive reinforcement*, a pleasant stimulus (often, a food reward) is presented following the behaviour. This increases the behaviour's likelihood. In *negative reinforcement*, an unwanted stimulus (often, as here, a painful stimulus) is removed following the behaviour, also increasing the behaviour. In *positive punishment*, an unwanted stimulus is presented following the behaviour, decreasing the behaviour's likelihood in the future.

China et al. very briefly mention the use of e-collars as a means of delivering punishment for unwanted behaviours, such as worrying livestock or killing wildlife, in their introduction, but the rest of their paper focuses entirely on using e-collars to deliver negative reinforcement to increase a desired behaviour. That is, in Group E the behaviours of returning to the trainer and sitting on command are negatively reinforced by the removal of an unpleasant stimulus, shock. I too think this is not a preferred training method, and am surprised that it worked as well as it did – well enough that it largely overcame the greatly reduced positive reinforcement for Group E compared with Group C2.

[6] There is no implied value judgment in the terms 'positive' and 'negative'. They are purely mathematical – is a stimulus being added or subtracted?

Consider recall behaviour. The aim of the training is that the dog will return to the trainer when summoned. Training this behaviour using positive reinforcement means that the reinforcer (food reward) is delivered when the behaviour is complete. Earlier parts of the behavioural chain, or sequence of behaviours, will be reinforced by increasing spatial and temporal proximity to the food reward, as well as by conditioned reinforcement – previously neutral stimuli, such as the trainer calling 'Good dog!', that have been paired with food in the past and so take on reinforcing properties. But training using negative reinforcement means that the reinforcer (cessation of the shock) comes only at the beginning of the behavioural chain – i.e., as soon as the dog starts to return. This is far less likely to be effective, not because of the nature of the reinforcer but because of the time at which it occurs during the behavioural chain. An analogy would be paying a plumber as soon as they arrive at your house, rather than when the work is completed – this is unlikely to be effective.

Given these considerations, the lack of a substantial difference between the results of Groups E and C2 is surprising. My tentative interpretation is that obedience to the command 'Come' was already well established before any training took place. This is consistent with both the lack of clear changes in behaviour for any group, no matter how they were trained, and the satisfactorily short latencies to respond for all groups even at the beginning of training.

Punishing unwanted behaviours:

Because of their (misguided, I think) focus on e-collars as a way to deliver negative reinforcement for a desired behaviour, China et al. almost entirely ignore the main indication for using an e-collar. This is to deliver punishment for an unwanted behaviour. In dogs, such behaviours may be relatively minor, such as continuous barking, voiding indoors, digging or damaging property. But they may be more serious, such as attacking livestock or wildlife, or aggression towards other dogs or humans. Without intervention, behaviours such as these are likely to result in destruction of the dog. While reinforcing alternative behaviours should certainly be part of any intervention, direct punishment of the unwanted behaviour as well is very likely to be needed to produce the necessary change.[7]

E-collars are likely to be a particular effective way to punish unwanted behaviours for several reasons. Firstly, as well as the punishment mechanism already discussed, Pavlovian

[7] In all species' natural environments, learning involves both positive reinforcement and punishment. We learn, for example, to drive within the speed limit because not doing so is punished. We are seldom positively reinforced for driving legally – we do so to avoid punishment.

conditioning processes also operate. If an electric shock is delivered when a dog attacks a sheep, for example, the sheep will become a Pavlovian negative conditioned stimulus. That is, the association of the sheep with shock will result in the sheep taking on (for the dog) some of the aversive properties of the shock. This will lead the dog to avoid the sheep in future, just as it would avoid the shock. Secondly, the role of the trainer/owner in delivering the shock is not obvious to the dog and, with a remote control, may be entirely covert. This means that the dog will associate the sheep, not the owner, with the shock. The owner will therefore not become aversive to the dog by the Pavlovian mechanism described above, as would be the case with other forms of punishment, such as smacking the dog's muzzle. Thirdly, the punishing shock can be delivered as soon as the dog initiates the unwanted behaviour, and at a distance. Finally, as in the use of buried electronic 'fences' that can operate an e-collar when the dog approaches the fence, they can be used to discourage straying without the need to restrain the dog physically and in the owner's absence.

There is a well-established literature on the efficacy of e-collars at reducing serious undesirable behaviour by dogs, which I will not rehearse here[8]. The point of this discussion is that China et al. do not address the main reason for e-collar use at all. Their paper simply attempts to show that e-collars do not confer an advantage when used to deliver negative reinforcement for two behaviours that are particularly unsuited to e-collar training.

Conclusions:

China et al. purport to show that training including e-collars was not more effective than training using only positive reinforcement at increasing dogs' compliance with the commands 'Come' and 'Sit'. There are sufficient problems with their procedures, research design and data analysis to make even this limited conclusion unreliable, and it is noteworthy that each of these problems can be expected to have an effect in the direction that would not support e-collar use. This is a biased comparison. However, while disagreeing that they have provided empirical support for their claim, my *a priori* analysis of the behavioural principles operating leads me to agree that using e-collars as a way to arrange negative reinforcement for a desired

[8] I have contributed to this literature myself. Most relevant to the discussion here, we have shown that a single pairing of e-collar shock with stimuli associated with the endangered New Zealand kiwi, such as nesting material, kiwi faeces, a silhouette of a kiwi and a stuffed kiwi, will lead to reliable avoidance of those stimuli in the future. This has the potential to allow dogs that have undergone such training to be licensed to enter kiwi habitat for purposes such as pig hunting (wild pigs are in themselves a threat to kiwi), benefiting kiwi, and dogs, and the dogs' owners. Such a result could certainly not be achieved by positive reinforcement – 'reward-based learning' – alone.

response is problematic. I too would not recommend that as the preferred method, although this research leaves open the possibility that negative reinforcement combined with an equal density of positive reinforcement may confer an advantage, even when training a behaviour that is *a priori* unsuited to e-collar use.

More seriously, China et al. overgeneralize their claimed result to suggest, at least by implication, that e-collars should never be used to train dogs. In doing so, they entirely disregard the most obvious indication for e-collar use, discouragement of an unwanted behaviour. Their research simply does not address this application of e-collars, even though e-collars are *a priori* much more suited to it than to the training of a desired behaviour, and there is a substantial literature supporting their effectiveness in such situations, particularly when the unwanted behaviour is serious and the owner is either at a distance from the dog or absent entirely. China et al.'s conclusion is not justified.

There may be ethical reasons that a government may decide to ban e-collars, and I am silent on that issue. My purpose in providing this commentary is to dismiss the argument that China et al. have provided empirical evidence that e-collars are ineffective. A policy decision that relies heavily on China et al. would not be a decision based on evidence.

Author's qualifications to comment:

Douglas Elliffe is Professor of Psychology at The University of Auckland, New Zealand. He is a former Head of the School of Psychology and currently Deputy Dean of the Faculty of Science. His main research area is behavioural psychology, both in animals and humans. He is a former Associate Editor of the flagship *Journal of the Experimental Behavior*, and regularly reviews manuscripts for a wide variety of behavioural journals. His 79 published papers include research, relevant to this commentary, on problem behaviours in dogs, on efficacy of e-collars to discourage predation, unwanted vocalization, and property damage by dogs, and on inferential statistics, as well as a large body of basic research on the principles of reinforcement and punishment, and how behaviour is controlled by its consequences.

This commentary is provided on a *pro bono* basis to Abzed Political and Media Relations in partial fulfilment of the critic and conscience obligations of Universities and their academic staff. There is no financial relationship between Abzed and Douglas Elliffe.