
Nearest Neighbour with Bandit Feedback

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 In this paper we adapt the nearest neighbour rule to the contextual bandit problem.
2 Our algorithm handles the fully adversarial setting in which no assumptions at all
3 are made about the data-generation process. When combined with a sufficiently
4 fast data-structure for (perhaps approximate) adaptive nearest neighbour search,
5 such as a navigating net, our algorithm is extremely efficient - having a per trial
6 running time polylogarithmic in both the number of trials and actions, and taking
7 only quasi-linear space. We give generic regret bounds for our algorithm and
8 further analyse them in a semi-stochastic setting. A side result of this paper is
9 that, when applied to the online classification problem with stochastic labels, our
10 algorithm can have sublinear regret whilst only finding a single nearest neighbour
11 per trial - in stark contrast to the k -nearest neighbours algorithm.

12 1 Introduction

13 In this paper we adapt the classic *nearest neighbour* rule to the contextual bandit problem and develop
14 an extremely efficient algorithm. The problem proceeds in trials, where on trial t : (1) a *context* x_t is
15 revealed to us, (2) we must select an *action* a_t , and (3) the *loss* $\ell_{t,a_t} \in [0, 1]$ of action a_t on trial t is
16 revealed to us. We assume that the contexts are points in a metric space and the distance between two
17 contexts represents their similarity. A *policy* is a mapping from contexts to actions and the inductive
18 bias of our algorithm is towards learning policies that typically map similar contexts to similar actions.
19 Our main result has absolutely no assumptions whatsoever about the generation of the context/loss
20 sequence and has no restriction on what policies we can compare our algorithm to.

21 Our algorithm requires, as a subroutine, a data-structure that performs c -nearest neighbour search.
22 This data-structure must be *adaptive* - in that new contexts can be inserted into it over time. An
23 example of such a data-structure is the *Navigating net* [15] which, given mild conditions on our
24 metric and dataset, performs both search and insertion in polylogarithmic time. When utilising a
25 data-structure of this speed our algorithm is extremely efficient - with a per-trial time complexity
26 polylogarithmic in both the number of trials and actions, and requiring only quasi-linear space.

27 As an example we will further analyse the special case of the contextual bandit problem in which the
28 context sequence is drawn i.i.d. from a probability distribution over the d -dimensional hypercube,
29 whilst the loss vectors can still be arbitrary. In this case, for any policy y with a finite-volume decision
30 boundary, our algorithm achieves sub-linear regret w.r.t. y without the need to know any parameters.

31 A side result of this paper is that, when applied to the online classification task with stochastic labels,
32 our algorithm can achieve sublinear regret whilst only finding *one* nearest neighbour per trial: in stark
33 contrast to the k -nearest neighbour algorithm. Our algorithm can hence be viewed also as a potentially
34 faster alternative to k -nearest neighbours when faced with the online classification problem.

35 In the course of this paper we develop some novel algorithmic techniques, including a new algorithmic
36 framework CANPROP and efficient algorithms for searching in trees, which may find further
37 application.

38 We now describe related works. The bandit problem [17] was first introduced in [22] but was
39 originally studied in the stochastic setting in which all losses are drawn i.i.d. at random [16], [1],
40 [2]. However, our world is very often not i.i.d. stochastic. The work of [3] introduced the seminal
41 EXP3 algorithm which handled the case in which the losses were selected arbitrarily. This work also
42 introduced the EXP4 algorithm for contextual bandits. In general this algorithm is exponential time
43 but in some situations can be implemented in polynomial time - such as their EXP3.S algorithm,
44 which was a bandit version of the classic FIXEDSHARE algorithm [13]. In [11] the EXP3.S setting
45 was greatly generalised to the situation in which the contexts were vertices of a graph. They utilised
46 the methodology of [7], [14] and [12] in order to develop extremely efficient algorithms. Although
47 inspiring this work, these algorithms cannot be utilised in our situation as they inherently require
48 the set of queried contexts to be known a-priori. In the stochastic case another class of contextual
49 bandit problems are *linear bandits* [18], [4] in which the contexts are mappings from the actions
50 into \mathbb{R}^d . Here the queried contexts need not be known in advance but the losses must be drawn i.i.d.
51 from a distribution that has mean linear in the respective context. The k nearest neighbour algorithm
52 was first analysed in [5]. The work [21] utilised the k nearest neighbour methodology and the works
53 [8] and [16] to handle a stochastic contextual bandit problem. However, their setting is extremely
54 more restricted than ours. In particular, the context/loss pairs must be drawn i.i.d. at random and the
55 probability distribution they are sampled from must obey certain strict conditions. In addition, on
56 each trial the contexts seen so far must be ordered in increasing distance from the current context
57 and operations must be performed on this sequence, making their algorithm exponentially slower
58 than ours. Our algorithm utilises the works of [19] and [6] as subroutines. It should be noted that the
59 later work, which was based on [20], was improved on in [10] - we leave it as an open problem as to
60 whether we can utilise their work in our algorithm.

61 2 Notation

62 Let \mathbb{N} be the set of natural numbers not including 0. Given a natural number $m \in \mathbb{N}$ we define
63 $\llbracket m \rrbracket := \{j \in \mathbb{N} \mid j \leq m\}$. Given a predicate p we define $\llbracket p \rrbracket := 1$ if p is true and $\llbracket p \rrbracket := 0$ otherwise.
64 We define $\log(\cdot)$ and $\ln(\cdot)$ to be the logarithms with base 2 and e respectively. Given sets \mathcal{A} and \mathcal{B}
65 we denote by $\mathcal{B}^{\mathcal{A}}$ the set of all functions $f : \mathcal{A} \rightarrow \mathcal{B}$ and by $2^{\mathcal{A}}$ the set of all subsets of \mathcal{A} .

66 All trees in this paper are considered rooted. Given a tree \mathcal{J} we denote its root by $r(\mathcal{J})$, its vertex set
67 by \mathcal{J} , its leaves by \mathcal{J}^* , and its internal vertices by \mathcal{J}^\dagger . Given a vertex v in a tree \mathcal{J} we denote its
68 parent by $\uparrow_{\mathcal{J}}(v)$ and the subtree of all its descendants by $\downarrow_{\mathcal{J}}(v)$. Given an internal node v in a (full)
69 binary tree \mathcal{J} we denote its left and right children by $\triangleleft_{\mathcal{J}}(v)$ and $\triangleright_{\mathcal{J}}(v)$ respectively. Internal nodes
70 v in a (full) ternary tree \mathcal{J} have an additional child $\nabla_{\mathcal{J}}(v)$ called the *centre* child. Given vertices
71 v and v' in a tree \mathcal{J} we denote by $\Gamma_{\mathcal{J}}(v, v')$ the *least common ancestor* of v and v' : i.e. the vertex
72 of maximum depth which is an ancestor of both v and v' . We will drop the subscript \mathcal{J} in all these
73 functions when unambiguous. Given a tree \mathcal{J} , a *subtree* of \mathcal{J} is a tree whose edge set is a subset of
74 that of \mathcal{J} .

75 Given a probability distribution μ we write $x \sim \mu$ to signify that x is a random element drawn from μ .
76 Given, in addition, some $m \in \mathbb{N}$, we define μ^m to be the probability distribution over sets formed by
77 drawing m elements i.i.d. with replacement from distribution μ . With a slight overloading of notation
78 we denote the uniform distribution over $[0, 1]$ also by $[0, 1]$.

79 3 Problem and Result

80 3.1 The Contextual Bandit Problem

81 We consider the following game between *Learner* (us) and *Nature* (our adversary). We have K actions
82 and a metric space (\mathcal{C}, Δ) where \mathcal{C} is a (possibly infinite) set of *contexts* and for all $x, x' \in \mathcal{C}$ we have
83 that $\Delta(x, x')$ is the *distance* from x to x' . Learning proceeds in T trials. A-priori Nature chooses a
84 sequence of contexts $\mathcal{X} = \{x_t \mid t \in [T]\} \subseteq \mathcal{C}$ and a sequence of loss vectors $\{\ell_t \mid t \in [T]\} \subseteq [0, 1]^K$,
85 but does not reveal them to Learner. On the t -th trial the following happens:

- 86 1. Nature reveals x_t to Learner.
- 87 2. Learner chooses some action $a_t \in [K]$.
- 88 3. Nature reveals ℓ_{t, a_t} to Learner.

89 A *policy* is a function from \mathcal{C} into $[K]$. i.e. a policy associates each possible context with an action.
 90 Given a policy $y : \mathcal{C} \rightarrow [K]$ we define the *y-regret* of Learner as:

$$R(y) := \sum_{t \in [T]} \ell_{t, a_t} - \sum_{t \in [T]} \ell_{t, y(x_t)}$$

91 which is the difference between the total cumulative loss suffered by Learner and that which Learner
 92 would have suffered if it had instead chosen a_t equal to $y(x_t)$ for all trials t .

93 3.2 The (k) Nearest Neighbour Classifier

94 We now digress from the contextual bandit problem in order to study the nearest neighbour methodol-
 95 ogy. The *nearest neighbour classifier* is a method to solve the following supervised learning problem.
 96 We assume that there exists an unknown function $y : \mathcal{C} \rightarrow [K]$. We are given a finite set $\mathcal{S} \subseteq \mathcal{C}$ along
 97 with the restriction of y onto \mathcal{S} . We are then asked to predict the value of $y(x)$ for some given $x \in \mathcal{C}$.
 98 The nearest neighbour classifier works by first finding the nearest neighbour \hat{x} of x , defined as:

$$\hat{x} := \operatorname{argmin}_{x' \in \mathcal{S}} \Delta(x, x'),$$

99 and then choosing $y(\hat{x})$ as its prediction of $y(x)$. In many important metric spaces the time taken to
 100 find such a nearest neighbour is in $\Theta(|\mathcal{S}|)$. This fact has led to the idea of instead using $y(\tilde{x})$ as our
 101 prediction, where $\tilde{x} \in \mathcal{S}$ is an arbitrary *c-nearest neighbour* which is defined as satisfying:

$$\Delta(x, \tilde{x}) \leq c \min_{x' \in \mathcal{S}} \Delta(x, x').$$

102 By utilising special data-structures the time complexity of finding, for any fixed $c > 1$, such a
 103 *c-nearest neighbour* is, for many metric spaces, only polylogarithmic in $|\mathcal{S}|$.

104 Given a probability distribution μ over \mathcal{C} , some $c \geq 1$ and some $m \in \mathbb{N}$ we define the *generalisation*
 105 *error* as:

$$g_m(\mu, y, c) := \mathbb{P}_{\mathcal{S} \sim \mu^m, x \sim \mu} \left[\exists z \in \mathcal{S} : \left(\Delta(x, z) \leq c \min_{x' \in \mathcal{S}} \Delta(x, x') \right) \wedge y(z) \neq y(x) \right]$$

106 which is the probability that it is possible for the nearest neighbour classifier to make a mistake on a
 107 context drawn from μ when \mathcal{S} is formed by drawing m contexts i.i.d. from μ .

108 We will now bound this quantity when in euclidean space. We first make the following definitions.
 109 For any $\delta > 0$ define the *δ -margin* of y by:

$$\mathcal{M}(y, \delta) := \{x \in \mathcal{C} \mid \exists x' \in \mathcal{C} : \Delta(x, x') \leq \delta \wedge y(x) \neq y(x')\} \quad (1)$$

110 which is the set of contexts that are at distance no more than δ from the *decision boundary* of y . The
 111 *volume* (w.r.t. μ) of the decision boundary is then given by:

$$\alpha(y, \mu) := \lim_{\delta \rightarrow 0} \frac{\mu(\mathcal{M}(y, \delta))}{\delta}. \quad (2)$$

112 When in euclidean space the following theorem bounds the generalisation error:

113 **Theorem 3.1.** *Given $\mathcal{C} := [0, 1]^d$ and Δ is the euclidean metric then for any $y : \mathcal{C} \rightarrow [K]$, $c \geq 1$,
 114 $\epsilon > 0$, and probability distribution μ such that the probability density of μ is always at least ϵ , we
 115 have:*

$$g_m(\mu, y, c) \in \tilde{\mathcal{O}} \left(c \alpha(y, \mu) (\epsilon m)^{-1/d} \right).$$

116 So far we have only considered deterministic functions $y : \mathcal{C} \rightarrow [K]$ with decision boundaries of finite
 117 volume. But what happens if instead we have that $y(x)$ is drawn from some probability distribution
 118 dependent on x (which is itself drawn from μ). Here, the *Bayes optimal classifier* is defined as that
 119 which always predicts $y^*(x) := \operatorname{argmax}_{a \in [K]} \mathbb{P}[y(x) = a \mid x]$ as the label of x . In general, even if
 120 $g_m(\mu, y^*, c) \in o(1)$, the probability of making a mistake with the nearest neighbour classifier does
 121 not approach that of the Bayes optimal classifier as $m \rightarrow \infty$. In order to converge optimally, the
 122 *k-nearest neighbour classifier* was introduced. In this algorithm, when given a context $x \in \mathcal{C}$, the *k*
 123 nearest neighbours to x are found and the predicted value of $y(x)$ is decided by majority vote. In
 124 order to converge optimally we need that $k \rightarrow \infty$ as $m \rightarrow \infty$.

125 A remarkable side-result of this paper is that, given $g_m(\mu, y^*, c) \in \mathcal{O}(m^{-p})$ for some $p > 0$, our
 126 algorithm can be applied to learning this situation online whilst only finding *one* nearest neighbour
 127 per trial. Since the additional overhead of our algorithm is so small it can be significantly faster than
 128 *k-nearest neighbours*. We strongly suspect that we don't need the condition on $g_m(\mu, y^*, c)$ if we are
 129 working in a bounded subset of euclidean space and $\mathbb{P}[y(x) = a \mid x]$ is Lipschitz.

130 3.3 Adaptive Nearest Neighbour Search

131 Our algorithm will require a data-structure for performing *adaptive nearest neighbour search*. This
 132 problem is as follows. We maintain a finite set $\mathcal{S} \subseteq \mathcal{C}$. At any point in time we must either (1)
 133 insert a new context into the set \mathcal{S} and update the data-structure, or (2) given a context, utilise the
 134 data-structure to find a c -nearest neighbour in the set \mathcal{S} .

135 We are especially interested in data-structures that can do both operations in a time polylogarithmic
 136 in $|\mathcal{S}|$. An example of such a data-structure is the *navigating net* [15] which has time complexity
 137 $\tilde{\mathcal{O}}(\ln(|\mathcal{S}|))$ given that $c > 1$, the set $|\mathcal{S}|$ is of bounded doubling dimension (w.r.t. Δ) and has aspect
 138 ratio (ratio between the largest and smallest distances between contexts in \mathcal{S}) polynomial in $|\mathcal{S}|$, as is
 139 the case in many applications and can be enforced by quantisation when working in a bounded subset
 140 of euclidean space. We note that the $\tilde{\mathcal{O}}$ hides a constant factor that is exponential in the doubling
 141 dimension of \mathcal{S} . In high-dimensional applications our dataset will often lie on a low-dimensional
 142 manifold so this factor should be small.

143 3.4 Our Results

144 We now turn back to the contextual bandit problem and give our main results.

145 **Theorem 3.2.** *Consider the contextual bandit problem defined in Section 3.1. Suppose that for all*
 146 *trials $t > 1$ we are given, in addition to x_t , a context $n(x_t)$ which satisfies:*

$$n(x_t) \in \{x_s \mid s \in [t-1]\}.$$

147 *Our algorithm CBNN takes a single parameter $\rho > 0$ and, for all policies $y : \mathcal{C} \rightarrow [K]$ simultane-*
 148 *ously, obtains an expected y -regret bounded by:*

$$\mathbb{E}[R(y)] \in \tilde{\mathcal{O}} \left(\left(\rho + \frac{\Phi(y)}{\rho} \right) \sqrt{KT} \right)$$

149 *where:*

$$\Phi(y) := 1 + \sum_{t \in [T] \setminus \{1\}} \mathbb{I}[y(x_t) \neq y(n(x_t))]$$

150 *and the expectation is taken over the randomisation of the algorithm. CBNN needs no initialisation*
 151 *time and has a per-trial time complexity of:*

$$\mathcal{O}(\ln(T)^2 \ln(K)).$$

152 We note that, although n can be any valid function, we are particularly interested in the case that
 153 $n(x_t)$ is a c -nearest neighbour of x_t , i.e. that we have:

$$\Delta(x_t, n(x_t)) \leq c \min_{s \in [t-1]} \Delta(x_t, x_s). \quad (3)$$

154 In this case finding $n(x_t)$ typically requires only $\tilde{\mathcal{O}}(\ln(T))$ time per trial when using a navigating
 155 net or other fast data-structure for adaptive nearest neighbour search, as explained in Section 3.3.
 156 Furthermore, the quantity $\Phi(y)$ can now be bounded in a way that is dependent only on the set of
 157 queried contexts \mathcal{X} and not their order. This bound is given in the following theorem.

158 **Theorem 3.3.** *Suppose we have a policy $y : \mathcal{C} \rightarrow [K]$. For any context $x \in \mathcal{C}$ we define $\gamma(x, y) :=$
 159 $\max\{\delta \geq 0 \mid x \notin \mathcal{M}(y, \delta)\}$ which is the distance of x from the decision boundary of y . Then when
 160 Equation (3) is satisfied we have that $\Phi(y)$ is no greater than the minimum cardinality of any set
 161 $\mathcal{S} \subseteq \mathcal{C}$ in which for all $t \in [T]$ there exists $x \in \mathcal{S}$ with $\Delta(x, x_t) < \gamma(x, y)/3c$.*

162 A direct corollary of this theorem is that for all $\delta > 0$ we have that:

$$\Phi(y) \leq N_\delta(\mathcal{X}) + |\mathcal{X} \cap \mathcal{M}(4c\delta, y)|$$

163 where $N_\delta(\mathcal{X})$ is the (external) covering number of \mathcal{X} with radius δ , and $|\mathcal{X} \cap \mathcal{M}(4c\delta, y)|$ is the
 164 number of contexts in \mathcal{X} lying within distance $4c\delta$ of the decision boundary.

165 It will be common in applications that the set \mathcal{X} of observed contexts will come from a finite set
 166 of separated clusters and there will be a good policy y which, on any such cluster, is constant on
 167 that cluster. Theorem 3.3 then implies that, as long as the inter-cluster distances are positive and the

168 clusters have finite covering numbers (which is guaranteed in many metric spaces), then $\Phi(y)$ will be
 169 bounded independent of T and hence, by Theorem 3.2, the y -regret of CBNN will scale like $\tilde{O}(\sqrt{T})$,
 170 whatever the value of ρ .

171 However, it will not always be the case that the dataset splits into such clusters. We shall investigate
 172 what happens when this is not the case by restricting ourselves to the situation in which the contexts
 173 $\{x_t \mid t \in [T]\}$ are drawn i.i.d. from a probability distribution μ . Here we have, by linearity of
 174 expectation, that:

$$\mathbb{E}[\Phi(y)] \leq 1 + \sum_{t \in [T]} g_t(\mu, y, c).$$

175 When in euclidean space, theorems 3.1 and 3.2 then lead to the following theorem:

176 **Theorem 3.4.** *Consider the contextual bandit problem defined in Section 3.1. Suppose that $\mathcal{C} =$
 177 $[0, 1]^d$, Δ is the euclidean metric, and the contexts are drawn i.i.d. at random from a probability
 178 distribution μ with density always at least $\epsilon > 0$. Note that the loss vectors can be arbitrary. Set ρ
 179 equal to $T^{-(d-1)/d} c^{-1/2}$. Then when Equation (3) is satisfied we have, for all policies $y : \mathcal{C} \rightarrow [K]$
 180 simultaneously, that the y -regret of CBNN is bounded by:*

$$\mathbb{E}[R(y)] \in \tilde{O} \left((\epsilon^{-1/d} \alpha(y, \mu) + 1) c^{1/2} K^{1/2} T^{(2d-1)/(2d)} \right)$$

181 where $\alpha(y, \mu)$ is the volume (w.r.t. μ) of the decision boundary of y as defined in equations (1) and
 182 (2). The existence of such an ϵ can be relaxed (with an effect on the bound) but we assume it for
 183 simplicity.

184 Note that, given the decision boundary of y is of finite volume, the expected regret is guaranteed to
 185 be sub-linear in T . This implies that the per-trial performance of CBNN approaches that of always
 186 selecting $a_t = y(x_t)$. We note that if T is unknown or infinite (i.e. learning never stops) then a
 187 simple doubling trick can be used to make the algorithm parameter-free (with no knowledge of μ).
 188 The fact that, in this non-separated case, the regret scales like $\tilde{O}(T^{(2d-1)/(2d)})$ is a facet of the well
 189 known *curse of dimensionality* and is the price we pay for being able to learn from such a vast class of
 190 policies. We note that in many high-dimensional applications the set \mathcal{X} will lie on a low-dimensional
 191 manifold so that the curse of dimensionality will be significantly reduced.

192 4 The Algorithm

193 In this section we describe our algorithm CBNN and give the pseudocode for the novel subroutines.
 194 In appendices C to E we give a more detailed description of how CBNN works, and prove that it
 195 obtains its bounds.

196 To give the reader intuition, in Appendix B we describe our initial idea - an algorithm, based on EXP4
 197 [3] and *Belief propagation* [20], which attains our regret bound but is exponentially slower - taking a
 198 per-trial time of $\Theta(KT)$.

199 4.1 Cancellation Propagation

200 In this subsection we describe a novel algorithmic framework CANPROP for designing contextual
 201 bandit algorithms with a running time logarithmic in K . It is inspired by EXP3 [3], specialist algo-
 202 rithms [7] and online decision-tree pruning algorithms [9] but is certainly not a simple combination
 203 of these works. CBNN will be an efficient implementation of an instance of CANPROP. Although in
 204 general CANPROP requires a-priori knowledge of the set $\mathcal{X} := \{x_t \mid t \in [T]\}$, CBNN is designed in
 205 a way that, crucially, does not need this set to be known.

206 We assume, without loss of generality, that K and T are integer powers of two. CANPROP, which
 207 takes a parameter $\eta > 0$, works on a full, balanced binary tree \mathcal{B} with leaves $\mathcal{B}^* = [K]$. On every
 208 trial t each pair $(v, \mathcal{S}) \in \mathcal{B} \times 2^{\mathcal{X}}$ has a weight $w_t(v, \mathcal{S}) \in [0, 1]$. These weights induce a function
 209 $\theta_t : \mathcal{B} \rightarrow [0, 1]$ defined by:

$$\theta_t(v) := \sum_{\mathcal{S} \in 2^{\mathcal{X}}} \mathbb{1}[x_t \in \mathcal{S}] w_t(v, \mathcal{S}).$$

210 On each trial t a root-to-leaf path $\{z_{t,j} \mid j \in [\log(K)] \cup \{0\}\}$ is sampled such that, given $z_{t,j}$, we
 211 have that $z_{t,(j+1)}$ is sampled from $\{\triangleleft(z_{t,j}), \triangleright(z_{t,j})\}$ with probability proportional to the value of θ_t

212 when applied to each of these vertices. The action a_t is then chosen equal to $z_{t,\log(K)}$. Once the loss
213 has been observed we climb back up the root-to-leaf path, updating the function w_t to w_{t+1} .
214 CANPROP (at trial t) is given in Algorithm 1. We note that if $w_{t+1}(v, \mathcal{S})$ is not set in the pseudocode
215 then it is defined to be equal to $w_t(v, \mathcal{S})$.

Algorithm 1 CANPROP at trial t

1: $v_{t,0} \leftarrow r(\mathcal{B})$ 2: for $j = 0, 1, \dots, (\log(K) - 1)$ do 3: for $v \in \{\triangleleft(v_{t,j}), \triangleright(v_{t,j})\}$ do 4: $\theta_t(v) \leftarrow \sum_{\mathcal{S} \in 2^{\mathcal{X}}} \mathbb{1}[x_t \in \mathcal{S}] w_t(v, \mathcal{S})$ 5: end for 6: $z_{t,j} \leftarrow \theta_t(\triangleleft(v_{t,j})) + \theta_t(\triangleright(v_{t,j}))$ 7: for $v \in \{\triangleleft(v_{t,j}), \triangleright(v_{t,j})\}$ do 8: $\pi_t(v) \leftarrow \theta_t(v) / z_{t,j}$ 9: end for 10: $\zeta_{t,j} \sim [0, 1]$ 11: if $\zeta_{t,j} \leq \pi_t(\triangleleft(v_{t,j}))$ then 12: $v_{t,j+1} \leftarrow \triangleleft(v_{t,j})$ 13: else 14: $v_{t,j+1} \leftarrow \triangleright(v_{t,j})$ 15: end if 16: end for	17: $a_t \leftarrow v_{t,\log(K)}$ 18: $\tilde{\pi}_t \leftarrow \prod_{j \in [\log(K)]} \pi_t(v_{t,j})$ 19: $\psi_{t,\log(K)} \leftarrow \exp(-\eta \ell_{t,a_t} / \tilde{\pi}_t)$ 20: for $j = \log(K), (\log(K) - 1), \dots, 1$ do 21: $\psi_{t,(j-1)} \leftarrow 1 - (1 - \psi_{t,j}) \pi_t(v_{t,j})$ 22: $\psi'_{t,j} \leftarrow \psi_{t,j} / \psi_{t,j-1}$ 23: if $v_{t,j} = \triangleleft(v_{t,j-1})$ then 24: $\tilde{v}_{t,j} \leftarrow \triangleright(v_{t,j-1})$ 25: else 26: $\tilde{v}_{t,j} \leftarrow \triangleleft(v_{t,j-1})$ 27: end if 28: for $\mathcal{S} \in 2^{\mathcal{X}} : x_t \in \mathcal{S}$ do 29: $w_{t+1}(v_{t,j}, \mathcal{S}) \leftarrow w_t(v_{t,j}, \mathcal{S}) \psi'_{t,j}$ 30: $w_{t+1}(\tilde{v}_{t,j}, \mathcal{S}) \leftarrow w_t(\tilde{v}_{t,j}, \mathcal{S}) / \psi_{t,j-1}$ 31: end for 32: end for
--	---

216 In Appendix C we give a general regret bound for CANPROP. For CBNN we set $\eta :=$
217 $\rho \sqrt{\ln(K) \ln(T) / KT}$ and for all $(v, \mathcal{S}) \in \mathcal{B} \times 2^{\mathcal{X}}$ set:

$$w_1(v, \mathcal{S}) := \frac{1}{4 \log(T)} \sum_{i \in [\log(T)]} \prod_{x \in \mathcal{X} \setminus \{x_1\}} \left(\sigma(x, \mathcal{S}) \frac{2^i}{T} + (1 - \sigma(x, \mathcal{S})) \left(1 - \frac{2^i}{T} \right) \right) \quad (4)$$

218 where $\sigma(x, \mathcal{S}) := \mathbb{1}[\{x \in \mathcal{S}\} \neq \{n(x) \in \mathcal{S}\}]$. This choice gives us the regret bound in Theorem 3.2.
219 We note that CBNN will be implemented in such a way that \mathcal{X} and n need not be known a-priori.

220 4.2 Ternary Search Trees

221 As we shall see, CBNN works by storing a binary tree $\mathcal{A}(v)$ at each vertex $v \in \mathcal{B}$. In order to perform
222 efficient operations on these trees we will utilise the rebalancing data-structure defined in [19] which
223 here we shall call a *ternary search tree* (TST) due to the fact that it is a generalisation of the classic
224 *binary search tree* and, as we shall show, has searching applications. However, as for binary search
225 trees, the applications of TSTs are more than just searching: we shall also utilise them for online
226 belief propagation.

227 We now define what is meant by a TST. Suppose we have a full binary tree \mathcal{J} . A TST of \mathcal{J} is a
228 full ternary tree \mathcal{D} which satisfies the following. The vertex set of \mathcal{D} is partitioned into two sets \mathcal{D}°
229 and \mathcal{D}^\bullet where each vertex $s \in \mathcal{D}$ is associated with a vertex $\mu(s) \in \mathcal{J}$ and every $s \in \mathcal{D}^\bullet$ is also
230 associated with a vertex $\mu'(s) \in \mathcal{J}$. In addition, each internal vertex $s \in \mathcal{D}^\dagger$ is associated with
231 a vertex $\xi(s) \in \mathcal{J}$. For all $u \in \mathcal{J}$ there exists a unique leaf $\Upsilon_{\mathcal{D}}(u) \in \mathcal{D}^*$ in which $\mu(\Upsilon_{\mathcal{D}}(u)) = u$.

232 For completeness we now describe the rules that a TST \mathcal{D} of \mathcal{J} must satisfy. We have that $r(\mathcal{D}) \in \mathcal{D}^\circ$
233 and $\mu(r(\mathcal{D})) := r(\mathcal{J})$. Each vertex $s \in \mathcal{D}$ represents a subtree $\hat{\mathcal{J}}(s)$ of \mathcal{J} . If $s \in \mathcal{D}^\circ$ then
234 $\hat{\mathcal{J}}(s) := \downarrow(\mu(s))$ and otherwise $\hat{\mathcal{J}}(s)$ is the set of descendants of $\mu(s)$ which are not proper
235 descendants of $\mu'(s)$. Given that $s \in \mathcal{D}^\dagger$ this subtree is *split* at the vertex $\xi(s)$ where if $s \in \mathcal{D}^\bullet$ we
236 have that $\xi(s)$ lies on the path from $\mu(s)$ to $\mu'(s)$. The children of s are then defined so that $\hat{\mathcal{J}}(\triangleleft(s)) =$
237 $\hat{\mathcal{J}}(s) \cap \downarrow(\triangleleft(\xi(s)))$ and $\hat{\mathcal{J}}(\triangleright(s)) = \hat{\mathcal{J}}(s) \cap \downarrow(\triangleright(\xi(s)))$ and $\hat{\mathcal{J}}(\nabla(s)) = \hat{\mathcal{J}}(s) \setminus (\hat{\mathcal{J}}(\triangleleft(s)) \cup \hat{\mathcal{J}}(\triangleright(s)))$,
238 i.e. $\xi(s)$ partitions the subtree $\hat{\mathcal{J}}(s)$ into the subtrees $\hat{\mathcal{J}}(\triangleleft(s))$, $\hat{\mathcal{J}}(\triangleright(s))$, and $\hat{\mathcal{J}}(\nabla(s))$. This process
239 continues recursively until $|\hat{\mathcal{J}}(s)| = 1$ in which case s is a leaf of \mathcal{D} .

240 For all binary trees \mathcal{J} in our algorithm we shall maintain a TST $\mathcal{H}(\mathcal{J})$ of \mathcal{J} with height $\mathcal{O}(\ln(|\mathcal{J}|))$.
 241 Such trees \mathcal{J} are *dynamic* in that on any trial it is possible that two vertices, u and u' , are added to
 242 the tree \mathcal{J} such that u' is inserted between a non-root vertex of \mathcal{J} and its parent, and u is designated
 243 as a child of u' . We define the subroutine $\text{REBALANCE}(\mathcal{H}(\mathcal{J}), u)$ as one which rebalances the TST
 244 $\mathcal{H}(\mathcal{J})$ after this insertion, so that the height of $\mathcal{H}(\mathcal{J})$ always remains in $\mathcal{O}(\ln(|\mathcal{J}|))$. The work of
 245 [19] describes how this subroutine can be implemented in a time of $\mathcal{O}(\ln(|\mathcal{J}|))$ and we refer the
 246 reader to this work for details (noting that they use different notation).

247 4.3 Contractions

248 At any trial t the contexts in $\{x_s \mid s \in [t]\}$ naturally form a tree by designating $n(x_s)$ as the parent
 249 of x_s . However, to utilise the TST data-structure we must only have binary trees. Hence, we will
 250 work with a (dynamic) full binary tree \mathcal{Z} which, on trial t , is a *binarisation* of the above tree. The
 251 relationship between these two trees is given by a map $\gamma : \mathcal{Z}_t \rightarrow \{x_s \mid s \in [t]\}$ where \mathcal{Z}_t is the tree
 252 \mathcal{Z} on trial t . For all $x \in \{x_s \mid s \in [t]\}$ we will always have a unique leaf $\tilde{\gamma}(x) \in \mathcal{Z}_t^*$ in which
 253 $\gamma(\tilde{\gamma}(x)) = x$. We also maintain a balanced TST $\mathcal{H}(\mathcal{Z})$ of \mathcal{Z} .

254 Algorithm 2 gives the subroutine GROW_t which updates \mathcal{Z} at the start of trial t . Note that GROW_t
 255 also defines a function $d : \mathcal{Z} \rightarrow \mathbb{N}$ such that $d(u)$ is the number of times the function n must be
 256 applied to $\gamma(u)$ to reach x_1 .

Algorithm 2 GROW_t which works on \mathcal{Z}

1: $u \leftarrow \tilde{\gamma}(n(x_t))$ 2: $u^* \leftarrow \uparrow(u)$ 3: $u' \leftarrow \text{NEWVERTEX}$ 4: $u'' \leftarrow \text{NEWVERTEX}$ 5: $\gamma(u') \leftarrow n(x_t)$ 6: $\gamma(u'') \leftarrow x_t$ 7: $\tilde{\gamma}(x_t) \leftarrow u''$ 8: if $u = \triangleleft(u^*)$ then 9: $\triangleleft(u^*) \leftarrow u'$	10: else 11: $\triangleright(u^*) \leftarrow u'$ 12: end if 13: $\triangleleft(u') \leftarrow u''$ 14: $\triangleright(u') \leftarrow u$ 15: $d(u') \leftarrow d(u)$ 16: $d(u'') \leftarrow d(u) + 1$ 17: $\text{REBALANCE}(\mathcal{H}(\mathcal{Z}), u'')$
---	--

257 A *contraction* (of \mathcal{Z}) is defined as a full binary tree \mathcal{J} in which the following holds. (1) The
 258 vertices of \mathcal{J} are a subset of those of \mathcal{Z} . (2) $r(\mathcal{J}) = r(\mathcal{Z})$. (3) Given a vertex $u \in \mathcal{J}$ we have
 259 $\triangleleft_{\mathcal{J}}(u) \in \downarrow_{\mathcal{Z}}(\triangleleft_{\mathcal{Z}}(u))$ and $\triangleright_{\mathcal{J}}(u) \in \downarrow_{\mathcal{Z}}(\triangleright_{\mathcal{Z}}(u))$. (4) Any leaf of \mathcal{J} is a leaf of \mathcal{Z} .

260 CBNN will maintain, on every vertex $v \in \mathcal{B}$, a contraction $\mathcal{A}(v)$ as well as a TST $\mathcal{H}(\mathcal{A}(v))$ of $\mathcal{A}(v)$.
 261 Given \mathcal{J} is one of these contractions, we also maintain, for all $i, i' \in \{0, 1\}$, all $u \in \mathcal{J}$ and all
 262 $j \in [\log(T)]$, a value $\tau_{i,i'}(\mathcal{J}, u, j) \in \mathbb{R}_+$. Technically these quantities, which depend on the above
 263 function d , define a sequence of *bayesian networks* on \mathcal{J} which is explained in Appendix D.3. For
 264 all $i \in \{0, 1\}$ and all $u \in \mathcal{J}$ we also maintain a value $\kappa_i(\mathcal{J}, u)$ initialised equal to 1.

265 On each of our contractions \mathcal{J} we will define, on trial t , a subroutine $\text{INSERT}_t(\mathcal{J})$ that simply
 266 modifies \mathcal{J} so that $\tilde{\gamma}(x_t)$ is added to its leaves. This subroutine is only called on certain trials t .
 267 Specifically, it is called on the contraction $\mathcal{A}(v)$ only when v is involved in CANPROP on trial t .
 268 Although the effect of this subroutine is simple to describe, its polylogarithmic-time implementation is
 269 quite complex. A function that is used many times during this subroutine is $\nu : \mathcal{Z} \times \mathcal{Z} \rightarrow \{\blacktriangleleft, \blacktriangleright, \blacktriangle\}$
 270 in which $\nu(u, u')$ is equal to $\blacktriangleleft, \blacktriangleright, \blacktriangle$ if u' is contained in $\downarrow_{\mathcal{Z}}(\triangleleft_{\mathcal{Z}}(u))$, in $\downarrow_{\mathcal{Z}}(\triangleright_{\mathcal{Z}}(u))$ or in neither,
 271 respectively. Algorithm 3 shows how to compute this function. Now that we have a subroutine for
 272 computing ν we can turn to the pseudocode for the subroutine $\text{INSERT}_t(\mathcal{J})$ in Algorithm 4. In the
 273 appendix we give a full description of how and why this subroutine works.

274 4.4 Online Belief Propagation

275 In this subsection we utilise the work of [6] in order to be able to efficiently compute the function θ_t
 276 that appears in CANPROP .

Algorithm 3 Computing $\nu(u, u')$ for $u, u' \in \mathcal{Z}$

```
1:  $\mathcal{E} \leftarrow \mathcal{H}(\mathcal{Z})$ 
2: if  $u = u'$  then
3:   return  $\blacktriangle$ 
4: end if
5:  $\tilde{s} \leftarrow \Upsilon_{\mathcal{E}}(u)$ 
6:  $\tilde{s}' \leftarrow \Upsilon_{\mathcal{E}}(u')$ 
7:  $s^* \leftarrow \Gamma_{\mathcal{E}}(\tilde{s}, \tilde{s}')$ 
8: for  $s \in \{\triangleleft(s^*), \nabla(s^*), \triangleright(s^*)\}$  do
9:   if  $\tilde{s} \in \Downarrow(s)$  then
10:     $\hat{s} \leftarrow s$ 
11:   end if
12:   if  $\tilde{s}' \in \Downarrow(s)$  then
13:     $\hat{s}' \leftarrow s$ 
14:   end if
15: end for
16: if  $\hat{s} \neq \nabla(s^*)$  then
17:   return  $\blacktriangle$ 
18: end if
19: if  $\xi(s^*) = u \wedge \hat{s}' = \triangleleft(s^*)$  then
20:   return  $\blacktriangleleft$ 
21: else if  $\xi(s^*) = u \wedge \hat{s}' = \triangleright(s^*)$  then
22:   return  $\blacktriangleright$ 
23: end if
24:  $s \leftarrow \hat{s}$ 
25: while TRUE do
26:   if  $s \in \mathcal{E}^\circ$  then
27:     return  $\blacktriangle$ 
28:   else if  $u = \xi(s) \wedge \triangleleft(s) \in \mathcal{E}^\bullet$  then
29:     return  $\blacktriangleleft$ 
30:   else if  $u = \xi(s) \wedge \triangleright(s) \in \mathcal{E}^\bullet$  then
31:     return  $\blacktriangleright$ 
32:   end if
33:   for  $s' \in \{\triangleleft(s), \nabla(s), \triangleright(s)\}$  do
34:     if  $\tilde{s} \in \Downarrow(s')$  then
35:        $s \leftarrow s'$ 
36:     end if
37:   end for
38: end while
```

Algorithm 4 The operation $\text{INSERT}_t(\mathcal{J})$ on a contraction \mathcal{J} of \mathcal{Z} at trial t

```
1:  $\mathcal{E} \leftarrow \mathcal{H}(\mathcal{Z})$ 
2:  $\mathcal{D} \leftarrow \mathcal{H}(\mathcal{J})$ 
3:  $s \leftarrow r(\mathcal{D})$ 
4:  $u_t \leftarrow \tilde{\gamma}(x_t)$ 
5: while  $s \in \mathcal{D}^\dagger$  do
6:   if  $\nu(\xi(s), u_t) = \blacktriangleleft$  then
7:      $s \leftarrow \triangleleft(s)$ 
8:   else if  $\nu(\xi(s), u_t) = \blacktriangleright$  then
9:      $s \leftarrow \triangleright(s)$ 
10:  else if  $\nu(\xi(s), u_t) = \blacktriangle$  then
11:     $s \leftarrow \nabla(s)$ 
12:  end if
13: end while
14:  $\hat{u} \leftarrow \mu(s)$ 
15:  $s \leftarrow r(\mathcal{E})$ 
16: while  $s \in \mathcal{E}^\dagger$  do
17:   if  $\nu(\xi(s), u_t) = \nu(\xi(s), \hat{u})$  then
18:     if  $\nu(\xi(s), u_t) = \blacktriangleleft$  then
19:        $s \leftarrow \triangleleft(s)$ 
20:     else if  $\nu(\xi(s), u_t) = \blacktriangleright$  then
21:        $s \leftarrow \triangleright(s)$ 
22:     else if  $\nu(\xi(s), u_t) = \blacktriangle$  then
23:        $s \leftarrow \nabla(s)$ 
24:     end if
25:   else
26:      $s \leftarrow \nabla(s)$ 
27:   end if
28: end while
29:  $u^* \leftarrow \mu(s)$ 
30:  $u' \leftarrow \uparrow_{\mathcal{J}}(\hat{u})$ 
31: if  $\hat{u} = \triangleleft_{\mathcal{J}}(u')$  then
32:    $\triangleleft_{\mathcal{J}}(u') \leftarrow u^*$ 
33: else
34:    $\triangleright_{\mathcal{J}}(u') \leftarrow u^*$ 
35: end if
36: if  $\nu(u^*, \hat{u}) = \blacktriangleleft$  then
37:    $\triangleleft_{\mathcal{J}}(u^*) \leftarrow \hat{u}$ 
38:    $\triangleright_{\mathcal{J}}(u^*) \leftarrow u_t$ 
39: else
40:    $\triangleright_{\mathcal{J}}(u^*) \leftarrow \hat{u}$ 
41:    $\triangleleft_{\mathcal{J}}(u^*) \leftarrow u_t$ 
42: end if
43: for  $i \in \{0, 1\}$  do
44:    $\kappa_i(\mathcal{J}, u^*) \leftarrow 1$ 
45:    $\kappa_i(\mathcal{J}, u_t) \leftarrow 1$ 
46: end for
47: for  $(j, i, i') \in [\log(T)] \times \{0, 1\} \times \{0, 1\}$  do
48:   for  $u \in \{u^*, \hat{u}, u_t\}$  do
49:      $\delta(u) \leftarrow d(u) - d(\uparrow_{\mathcal{J}}(u))$ 
50:     if  $i = i'$  then
51:        $\tau_{i, i'}(\mathcal{J}, u, j) \leftarrow 1 - \phi_{\delta(u)}(2^j/T)$ 
52:     else
53:        $\tau_{i, i'}(\mathcal{J}, u, j) \leftarrow \phi_{\delta(u)}(2^j/T)$ 
54:     end if
55:   end for
56: end for
57:  $\text{REBALANCE}(\mathcal{H}(\mathcal{J}), u_t)$ 
```

277 Given a vertex u in one of our contractions \mathcal{J} we define $\mathcal{F}(\mathcal{J}, u) := \{f \in \{0, 1\}^{\mathcal{J}} \mid f(u) = 1\}$ and
 278 then for all $j \in [\log(T)]$ define:

$$\Lambda(\mathcal{J}, u, j) := \prod_{f \in \mathcal{F}(\mathcal{J}, u)} \prod_{u' \in \mathcal{J} \setminus \{r(\mathcal{J})\}} \tau_{f(\uparrow_{\mathcal{J}}(u')), f(u')}(\mathcal{J}, u', j) \kappa_{f(u')}(\mathcal{J}, u').$$

279 As stated in the previous subsection, when a vertex $v \in \mathcal{B}$ becomes involved in CANPROP on trial t ,
 280 CBNN will add $\tilde{\gamma}(x_t)$ to the leaves of $\mathcal{A}(v)$ via the operation $\text{INSERT}_t(\mathcal{A}(v))$. In the appendix we
 281 shall show that for each such v we then have:

$$\theta_t(v) = \frac{1}{4 \log(T)} \sum_{j \in [\log(T)]} \Lambda(\mathcal{A}(v), \tilde{\gamma}(x_t), j).$$

282 We now outline how to compute this efficiently, deferring a full description for Appendix E.3. First
 283 note that for all contractions \mathcal{J} and all $(j, u) \in [\log(T)] \times \mathcal{J}$ we have that $\Lambda(\mathcal{J}, u, j)$ is of the exact
 284 form to be solved by the classic *Belief propagation* algorithm [20]. The work of [6] shows how to
 285 compute this term in logarithmic time by maintaining a data-structure based on a balanced TST of \mathcal{J} -
 286 in our case the TST $\mathcal{H}(\mathcal{J})$. Whenever, for some $i \in \{0, 1\}$ and $u' \in \mathcal{J}$, the value $\kappa_i(\mathcal{J}, u')$ changes,
 287 the data-structure is updated in logarithmic time.

288 We shall maintain, for each contraction \mathcal{J} , a set of $\log(T)$ such data-structures - one for each
 289 value of j . We define the subroutine $\text{EVIDENCE}(\mathcal{J}, u')$ as that which updates all these data-
 290 structures after $\kappa_i(\mathcal{J}, u')$ changes. We also make sure that the data-structures are updated whenever
 291 $\text{REBALANCE}(\mathcal{H}(\mathcal{J}), \cdot)$ is called. We then define the subroutine $\text{MARGINAL}(\mathcal{J}, u)$ as that which
 292 computes $\Lambda(\mathcal{J}, u, j)$ for each $j \in [\log(T)]$, and then sums the results and divides by $4 \log(T)$. Hence,
 293 the output of $\text{MARGINAL}(\mathcal{A}(v), \tilde{\gamma}(x_t))$ is equal to $\theta_t(v)$.

294 4.5 CBNN

295 Now that we have defined all our subroutines we give, in Algorithm 5, the algorithm CBNN which is
 296 an efficient implementation of CANPROP with initial weighting given in Equation (4).

Algorithm 5 CBNN at trial t

<pre> 1: GROW_t 2: $u_t \leftarrow \tilde{\gamma}(x_t)$ 3: $v_{t,0} \leftarrow r(\mathcal{B})$ 4: for $j = 0, 1, \dots, (\log(K) - 1)$ do 5: for $v \in \{\triangleleft(v_{t,j}), \triangleright(v_{t,j})\}$ do 6: $\text{INSERT}_t(\mathcal{A}(v))$ 7: $\theta_t(v) \leftarrow \text{MARGINAL}(\mathcal{A}(v), u_t)$ 8: end for 9: $z_{t,j} \leftarrow \theta_t(\triangleleft(v_{t,j})) + \theta_t(\triangleright(v_{t,j}))$ 10: for $v \in \{\triangleleft(v_{t,j}), \triangleright(v_{t,j})\}$ do 11: $\pi_t(v) \leftarrow \theta_t(v) / z_{t,j}$ 12: end for 13: $\zeta_{t,j} \sim [0, 1]$ 14: if $\zeta_{t,j} \leq \pi_t(\triangleleft(v_{t,j}))$ then 15: $v_{t,j+1} \leftarrow \triangleleft(v_{t,j})$ 16: else 17: $v_{t,j+1} \leftarrow \triangleright(v_{t,j})$ </pre>	<pre> 18: end if 19: end for 20: $a_t \leftarrow v_{t, \log(K)}$ 21: $\tilde{\pi}_t \leftarrow \prod_{j \in [\log(K)]} \pi_t(v_{t,j})$ 22: $\psi_{t, \log(K)} \leftarrow \exp(-\eta \ell_{t, a_t} / \tilde{\pi}_t)$ 23: for $j = \log(K), (\log(K) - 1), \dots, 1$ do 24: $\psi_{t, (j-1)} \leftarrow 1 - (1 - \psi_{t,j}) \pi_t(v_{t,j})$ 25: if $v_{t,j} = \triangleleft(v_{t, j-1})$ then 26: $\tilde{v}_{t,j} \leftarrow \triangleright(v_{t, j-1})$ 27: else 28: $\tilde{v}_{t,j} \leftarrow \triangleleft(v_{t, j-1})$ 29: end if 30: $\kappa_1(\mathcal{A}(v_{t,j}), u_t) \leftarrow \psi_{t,j} / \psi_{t, j-1}$ 31: $\kappa_1(\mathcal{A}(\tilde{v}_{t,j}), u_t) \leftarrow 1 / \psi_{t, j-1}$ 32: $\text{EVIDENCE}(\mathcal{A}(v_{t,j}), u_t)$ 33: $\text{EVIDENCE}(\mathcal{A}(\tilde{v}_{t,j}), u_t)$ 34: end for </pre>
--	--

297 5 Conclusion

298 In this paper we introduced the use of the nearest neighbour methodology for the fully adversarial
 299 contextual bandit problem when the contexts are selected from a metric space. We developed an
 300 extremely efficient algorithm CBNN. We gave a regret bound for CBNN and, as an example, further
 301 analysed it in the case in which the contexts (but not necessarily the losses) are drawn i.i.d. from a
 302 distribution on a multi-dimensional hypercube: where CBNN requires no knowledge of parameters.

303 **References**

- 304 [1] R. Agrawal. Sample mean based index policies by $o(\log n)$ regret for the multi-armed bandit
305 problem. *Advances in Applied Probability*, 27:1054 – 1078, 1995.
- 306 [2] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem.
307 *Machine Learning*, 47:235–256, 2002.
- 308 [3] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire. The nonstochastic multiarmed bandit
309 problem. *SIAM J. Comput.*, 32:48–77, 2002.
- 310 [4] W. Chu, L. Li, L. Reyzin, and R. E. Schapire. Contextual bandits with linear payoff functions.
311 In *International Conference on Artificial Intelligence and Statistics*, 2011.
- 312 [5] T. M. Cover and P. E. Hart. Nearest neighbor pattern classification. *IEEE Trans. Inf. Theory*,
313 13:21–27, 1967.
- 314 [6] A. L. Delcher, A. J. Grove, S. Kasif, and J. Pearl. Logarithmic-time updates and queries in
315 probabilistic networks. *J. Artif. Intell. Res.*, 4:37–59, 1995.
- 316 [7] Y. Freund, R. E. Schapire, Y. Singer, and M. K. Warmuth. Using and combining predictors that
317 specialize. In *Symposium on the Theory of Computing*, 1997.
- 318 [8] A. Garivier and O. Cappé. The kl-ucb algorithm for bounded stochastic bandits and beyond. In
319 *Annual Conference Computational Learning Theory*, 2011.
- 320 [9] D. P. Helmbold and R. E. Schapire. Predicting nearly as well as the best pruning of a decision
321 tree. *Machine Learning*, 27:51–68, 1995.
- 322 [10] M. Herbster, S. Pasteris, and F. Vitale. Online sum-product computation over trees. In *NIPS*,
323 2012.
- 324 [11] M. Herbster, S. Pasteris, F. Vitale, and M. Pontil. A gang of adversarial bandits. In *Neural*
325 *Information Processing Systems*, 2021.
- 326 [12] M. Herbster and J. Robinson. Online prediction of switching graph labelings with cluster
327 specialists. In *Neural Information Processing Systems*, 2018.
- 328 [13] M. Herbster and M. K. Warmuth. Tracking the best expert. *Machine Learning*, 32:151–178,
329 1995.
- 330 [14] W. M. Koolen, D. Adamskiy, and M. K. Warmuth. Putting bayes to sleep. In *NIPS*, 2012.
- 331 [15] R. Krauthgamer and J. R. Lee. Navigating nets: simple algorithms for proximity search. In
332 *ACM-SIAM Symposium on Discrete Algorithms*, 2004.
- 333 [16] T. L. Lai and H. E. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in*
334 *Applied Mathematics*, 6:4–22, 1985.
- 335 [17] T. Lattimore and C. Szepesvári. Bandit algorithms. 2020.
- 336 [18] L. Li, W. Chu, J. Langford, and R. E. Schapire. A contextual-bandit approach to personalized
337 news article recommendation. In *The Web Conference*, 2010.
- 338 [19] K. Matsuzaki and A. Morihata. Mathematical engineering technical reports balanced ternary-tree
339 representation of binary trees and balancing algorithms. 2008.
- 340 [20] J. Pearl. Reverend bayes on inference engines: A distributed hierarchical approach. *Probabilistic*
341 *and Causal Inference*, 1982.
- 342 [21] H. W. J. Reeve, J. C. Mellor, and G. Brown. The k-nearest neighbour ucb algorithm for
343 multi-armed bandits with covariates. *ArXiv*, abs/1803.00316, 2018.
- 344 [22] H. E. Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American*
345 *Mathematical Society*, 58:527–535, 1952.

346 A Guide to the Appendices

347 To give the reader some intuition behind CBNN we present, in Appendix B, our initial idea: an
 348 algorithm which obtains the regret bound of CBNN but is exponentially slower. We then give a
 349 detailed description of CBNN in appendices C to E. Specifically, in Appendix C we describe our
 350 novel algorithmic framework CANPROP. In Appendix D we describe contractions and bayesian
 351 networks on them, showing how CANPROP can be implemented with them. Finally, in Appendix E
 352 we describe TSTs and how they are used to perform our required operations efficiently. In Appendix
 353 F we prove, in order, all of the theorems stated in this paper.

354 B The Initial Idea

355 Here we describe our initial idea - an algorithm, based on EXP4 [3] and *Belief propagation* [20],
 356 which attains the regret bound of CBNN but is exponentially slower - taking a per-trial time of
 357 $\tilde{\Theta}(KT)$. Since this section is only to give intuition, and the results are surpassed by CBNN, we do
 358 not prove the statements made in this section.

359 To begin with we assume a-priori knowledge of the set $\mathcal{X} := \{x_t \mid t \in [T]\}$ and function n but the
 360 final algorithm will not need this knowledge. Without loss of generality assume that T is an integer
 361 power of two.

362 The algorithm is based on EXP4 [3] which we now describe. On every trial t we maintain a weighting
 363 $\hat{w}_t : [K]^{\mathcal{X}} \rightarrow [0, 1]$. We are free to choose any \hat{w}_1 satisfying:

$$\sum_{y \in [K]^{\mathcal{X}}} \hat{w}_1(y) = 1.$$

364 On each trial t the following happens:

- 365 1. x_t is revealed
- 366 2. For all $a \in [K]$ set $p_{t,a} \leftarrow \sum_{y \in [K]^{\mathcal{X}}} \mathbb{1}[y(x_t) = a] \hat{w}_t(y)$
- 367 3. Set $a_t \leftarrow a$ with probability proportional to $p_{t,a}$
- 368 4. Receive ℓ_{t,a_t}
- 369 5. For all $a \in [K]$ set $\hat{\ell}_{t,a} \leftarrow \mathbb{1}[a = a_t] \ell_{t,a_t} \|\mathbf{p}_t\|_1 / p_{t,a_t}$
- 370 6. For all $y \in [K]^{\mathcal{X}}$ set $\hat{w}_{t+1}(y) \leftarrow \hat{w}_t(y) \exp(-\eta \hat{\ell}_{t,y(x_t)})$

371 It is a classic result [3] that, for any policy $y : \mathcal{X} \rightarrow [K]$, the expected y -regret of EXP4 is bounded
 372 by:

$$\mathbb{E}[R(y)] \leq \frac{\eta KT}{2} - \frac{\ln(\hat{w}_1(y))}{\eta}.$$

373 If $i \in [\log(T)]$ is such that $2^i \leq \Phi(y) \leq 2^{i+1}$ then:

$$\ln \left(\left(\frac{2^i}{T(K-1)} \right)^{\Phi(y)} \left(1 - \frac{2^i}{T} \right)^{(T-1-\Phi(y))} \right) \in \mathcal{O} \left(\ln \left(\frac{KT}{|\Phi(y)|} \right) \Phi(y) \right)$$

374 so setting:

$$\eta := \rho \sqrt{\frac{1}{KT}}$$

375 and:

$$\hat{w}_1(y) := \frac{1}{K \log(T)} \sum_{i \in [\log(T)]} \left(\frac{2^i}{T(K-1)} \right)^{\Phi(y)} \left(1 - \frac{2^i}{T} \right)^{(T-1-\Phi(y))}$$

376 gives us our desired regret bound.

377 However, we have two issues - the algorithm takes exponential time and the set \mathcal{X} and function
 378 n need to be known a-priori. We will hence discuss how to bring the time complexity down to

379 $\tilde{\Theta}(KT)$ and with no a-priori knowledge. To do this first define, for all $i \in [\log(T)]$, the function
 380 $\hat{\tau}_i : [K] \times [K] \rightarrow [0, 1]$ by:

$$\hat{\tau}_i(a, a') := \llbracket a \neq a' \rrbracket \frac{2^i}{T(K-1)} + \llbracket a = a' \rrbracket \left(1 - \frac{2^i}{T}\right)$$

381 Note then that for all $y \in [K]^{\mathcal{X}}$ we have:

$$\hat{w}_1(y) \propto \sum_{i \in [\log(T)]} \left(\prod_{s \in [T] \setminus \{1\}} \hat{\tau}_i(y(x_s), y(n(x_s))) \right)$$

382 so that for all trials t :

$$\hat{w}_t(y) \propto \sum_{i \in [\log(T)]} \left(\prod_{s \in [T] \setminus \{1\}} \hat{\tau}_i(y(x_s), y(n(x_s))) \right) \prod_{s \in [T]} \exp(-\llbracket s < t \rrbracket \eta \hat{\ell}_{s, y(x_s)})$$

383 and hence for all $a \in [K]$ we have:

$$p_{t, a} \propto \sum_{i \in [\log(T)]} \sum_{y \in [K]^{\mathcal{X}}} \llbracket y(x_t) = a \rrbracket \left(\prod_{s \in [T] \setminus \{1\}} \hat{\tau}_i(y(x_s), y(n(x_s))) \right) \prod_{s \in [T]} \exp(-\llbracket s < t \rrbracket \eta \hat{\ell}_{s, y(x_s)})$$

384 For all $s \in [t]$ and $a \in [K]$ define:

$$\phi'_t(x_s, a) := \llbracket s < t \rrbracket \exp(-\eta \hat{\ell}_{s, a}) + \llbracket s = t \rrbracket$$

385 A crucial insight is that:

$$\sum_{y \in [K]^{\mathcal{X}}} \llbracket y(x_t) = a \rrbracket \left(\prod_{s \in [T] \setminus \{1\}} \hat{\tau}_i(y(x_s), y(n(x_s))) \right) \prod_{s \in [T]} \exp(-\llbracket s < t \rrbracket \eta \hat{\ell}_{s, y(x_s)})$$

386 is equal to:

$$\sum_{y \in [K]^{\mathcal{X}_t}} \llbracket y(x_t) = a \rrbracket \phi'_t(x_1, y(x_1)) \prod_{s \in [t] \setminus \{1\}} \hat{\tau}_s(y(x_s), y(n(x_s))) \phi'_t(x_s, y(x_s)) \quad (5)$$

387 which is why the algorithm needs only know $\{x_s \mid s \in [t]\}$ and $\{n(x_s) \mid s \in [t] \setminus \{1\}\}$. On trial t we
 388 construct a tree with vertex set $\{x_s \mid s \in [t]\}$ which is rooted at x_1 and is such that for all $s \in [t] \setminus \{1\}$
 389 we have that $n(x_s)$ is the parent of x_s . We note that computing the quantity in Equation (5) for all
 390 $a \in [K]$ can be done in a time of $\Theta(Kt)$ by Belief propagation [20] on this tree. Hence we have that
 391 p_t can be computed in a time of $\Theta(Kt \log(T))$ without a-priori knowledge of \mathcal{X} and n .

392 C Cancellation Propagation

393 We now turn to the description and analysis of our algorithm CBNN, starting with our novel
 394 algorithmic framework CANPROP.

395 C.1 The General CANPROP Algorithm

396 Let $\mathcal{X} := \{x_t \mid t \in [T]\}$. Note that we do not know \mathcal{X} a-priori but for now let's assume we do. We
 397 now introduce a general algorithmic framework CANPROP for handling contextual bandit problems
 398 with a per-trial time logarithmic in K . Without loss of generality assume that K is an integer power
 399 of two. Let \mathcal{B} be a full, balanced, oriented binary tree whose leaves are the set of actions $[K]$. Let
 400 $\mathcal{B}' := \mathcal{B} \setminus \{r(\mathcal{B})\}$. CANPROP takes a parameter $\eta \in \mathbb{R}_+$ called the *learning rate*. On each trial t
 401 CANPROP maintains a function:

$$w_t : \mathcal{B}' \times 2^{\mathcal{X}} \rightarrow [0, 1]$$

402 The function w_1 is free to be defined how one likes, as long as it satisfies the constraint that for all
 403 internal vertices $v \in \mathcal{B}^\dagger$ we have:

$$\sum_{\mathcal{S} \in 2^{\mathcal{X}}} (w_1(\triangleleft(v), \mathcal{S}) + w_1(\triangleright(v), \mathcal{S})) = 1$$

404 We now describe how CANPROP acts on trial t . For all $v \in \mathcal{B}'$ we define:

$$\theta_t(v) := \sum_{\mathcal{S} \in 2^{\mathcal{X}}} \mathbb{1}[x_t \in \mathcal{S}] w_t(v, \mathcal{S})$$

405 and for all $v \in \mathcal{B}^\dagger$ we define:

$$\pi_t(\triangleleft(v)) := \frac{\theta_t(\triangleleft(v))}{\theta_t(\triangleleft(v)) + \theta_t(\triangleright(v))} \quad ; \quad \pi_t(\triangleright(v)) := \frac{\theta_t(\triangleright(v))}{\theta_t(\triangleleft(v)) + \theta_t(\triangleright(v))}$$

406 As we shall see CANPROP needs only compute these values for $\mathcal{O}(\ln(K))$ vertices v . CANPROP
 407 samples a root-to-leaf path $\{v_{t,j} \mid j \in [\log(K)] \cup \{0\}\}$ as follows. $v_{t,0}$ is defined equal to $r(\mathcal{B})$. For
 408 all $j \in [\log(K) - 1] \cup \{0\}$, once $v_{t,j}$ has been sampled we sample $v_{t,(j+1)}$ from the probability
 409 distribution defined by:

$$\mathbb{P}[v_{t,(j+1)} = v] := \mathbb{1}[\uparrow(v) = v_{t,j}] \pi_t(v) \quad \forall v \in \mathcal{B}'$$

410 noting that $v_{t,(j+1)}$ is a child of $v_{t,j}$. We define:

$$\mathcal{P}_t := \{v_{t,j} \mid j \in [\log(K)] \cup \{0\}\}$$

411 CANPROP then selects:

$$a_t := v_{t,\log(K)}$$

412 and then receives the loss ℓ_{t,a_t} . The function w_t is then updated to w_{t+1} as follows. Firstly we define,

$$w_{t+1}(v, \mathcal{S}) := w_t(v, \mathcal{S}) \quad \forall (v, \mathcal{S}) \in \{v' \in \mathcal{B}' \mid \uparrow(v') \notin \mathcal{P}_t\} \times 2^{\mathcal{X}}$$

413 We then define:

$$\psi_{t,\log(K)} := \exp\left(\frac{-\eta \ell_{t,a_t}}{\prod_{j \in [\log(K)]} \pi_t(v_{t,j})}\right)$$

414 Once we have defined $\psi_{t,j}$ for some $j \in [\log(K)]$ we then define:

$$\psi_{t,(j-1)} := 1 - (1 - \psi_{t,j}) \pi_t(v_{t,j})$$

415

$$\beta_t(v) := \frac{\mathbb{1}[v \in \mathcal{P}_t] \psi_{t,j} + \mathbb{1}[v \notin \mathcal{P}_t]}{\psi_{t,(j-1)}} \quad \forall v \in \{\triangleleft(v_{t,(j-1)}), \triangleright(v_{t,(j-1)})\}$$

416

$$w_{t+1}(v, \mathcal{S}) := (\mathbb{1}[x_t \in \mathcal{S}] \beta_t(v) + \mathbb{1}[x_t \notin \mathcal{S}]) w_t(v, \mathcal{S}) \quad \forall (v, \mathcal{S}) \in \{\triangleleft(v_{t,(j-1)}), \triangleright(v_{t,(j-1)})\} \times 2^{\mathcal{X}}$$

417 The regret bound of CANPROP is given by the following theorem.

418 **Theorem C.1.** *Suppose we have a function $y : \mathcal{X} \rightarrow [K]$. For all $v \in \mathcal{B}$ define:*

$$\mathcal{Q}(v) := \{x \in \mathcal{X} \mid y(x) \in \downarrow(v)\}$$

419 *Then the expected y -regret of CANPROP is bounded by:*

$$\mathbb{E}[R(y)] \leq \frac{\eta K T}{2} - \frac{1}{\eta} \sum_{v \in \mathcal{B}'} \mathbb{1}[\mathcal{Q}(v) \neq \emptyset] \ln(w_1(v, \mathcal{Q}(v)))$$

420 C.2 Our Parameter Tuning

421 We now describe and analyse the initial weighting w_1 that we will use. Without loss of generality
 422 assume T is an integer power of two. Define $\mathcal{X}' := \mathcal{X} \setminus \{x_1\}$. For all $(x, \mathcal{S}) \in \mathcal{X}' \times 2^{\mathcal{X}}$ define:

$$\sigma(x, \mathcal{S}) := \mathbb{1}[\mathbb{1}[x \in \mathcal{S}] \neq \mathbb{1}[n(x) \in \mathcal{S}]]$$

423 For all $(v, \mathcal{S}) \in \mathcal{B}' \times 2^{\mathcal{X}}$ we define:

$$w_1(v, \mathcal{S}) := \frac{1}{4 \log(T)} \sum_{i \in [\log(T)]} \prod_{x \in \mathcal{X}'} \left(\sigma(x, \mathcal{S}) \frac{2^i}{T} + (1 - \sigma(x, \mathcal{S})) \left(1 - \frac{2^i}{T}\right) \right)$$

424 Given our parameter ρ we choose our learning rate as:

$$\eta := \rho \sqrt{\frac{\ln(K) \ln(T)}{K T}}$$

425 Given this initial weighting and learning rate, Theorem C.1 implies the following regret bound.

426 **Theorem C.2.** *Given w_1 and η are defined as above, then for any policy $y : \mathcal{C} \rightarrow [K]$ the expected
 427 y -regret of CANPROP is bounded by:*

$$\mathbb{E}[R(y)] \in \mathcal{O}\left(\left(\rho + \frac{\Phi(y)}{\rho}\right) \sqrt{\ln(K) \ln(T) K T}\right)$$

428 D Binarisation and Implementation with Contractions

429 D.1 A Sequence of Binary Trees

430 For any trial t we have a natural tree-structure on the set $\{x_{t'} \mid t' \in [t]\}$ formed by making $n(x_{t'})$ the
 431 parent of $x_{t'}$ for all $t' \in [t] \setminus \{1\}$. However, in order to utilise the methodology of [19] we need to work
 432 with binary trees. Hence, we now inductively define a sequence of binary trees $\{\mathcal{Z}_t \mid t \in [T] \setminus \{1\}\}$
 433 where the vertices of \mathcal{Z}_t are a subset of those of \mathcal{Z}_{t+1} . We also define a function $\gamma : \mathcal{Z}_T \rightarrow \mathcal{X}$. This
 434 function γ has the property that for any $t \in [T]$ and for any distinct leaves $u, u' \in \mathcal{Z}_t^*$ we have that
 435 $\gamma(u) \neq \gamma(u')$, and that:

$$\{\gamma(u) \mid u \in \mathcal{Z}_t^*\} = \{x_{t'} \mid t' \in [t]\}$$

436 We define \mathcal{Z}_2 to contain three vertices $\{r(\mathcal{Z}_2), \triangleleft(r(\mathcal{Z}_2)), \triangleright(r(\mathcal{Z}_2))\}$ where:

$$\gamma(r(\mathcal{Z}_2)) := \gamma(\triangleleft(r(\mathcal{Z}_2))) := x_1 \quad \text{and} \quad \gamma(\triangleright(r(\mathcal{Z}_2))) := x_2$$

437 Now consider a trial $t \in [T]$. We have that \mathcal{Z}_{t+1} is constructed from \mathcal{Z}_t via the following algorithm
 438 **GROW $_{t+1}$** :

- 439 1. Let u be the unique leaf in \mathcal{Z}_t^* in which $\gamma(u) = n(x_{t+1})$ and let $u^* := \uparrow(u)$.
- 440 2. Create two new vertices u' and u'' .
- 441 3. Set $\gamma(u') \leftarrow n(x_{t+1})$ and $\gamma(u'') \leftarrow x_{t+1}$.
- 442 4. If $u = \triangleleft(u^*)$ then set $\triangleleft(u^*) \leftarrow u'$. Else set $\triangleright(u^*) \leftarrow u'$.
- 443 5. Set $\triangleleft(u') \leftarrow u''$ and $\triangleright(u') \leftarrow u$

444 We also define a function $d : \mathcal{Z}_T \rightarrow \mathbb{N} \cup \{0\}$ as follows. Define $d'(x_1) := 0$ and for all $t \in [T] \setminus \{1\}$
 445 inductively define $d'(x_t) := d'(n(x_t)) + 1$. Finally define $d(u) := d'(\gamma(u))$ for all $u \in \mathcal{Z}_T$. Since
 446 for all $t \in [T]$ we have that the vertices of \mathcal{Z}_t are a subset of those of \mathcal{Z}_T we have that d also defines
 447 a function over \mathcal{Z}_t for all $t \in [T]$.

448 For all $t \in [T]$ we define u_t to be the unique leaf of \mathcal{Z}_t for which $\gamma(u_t) = x_t$.

449 D.2 Contractions

450 Our efficient implementation of CANPROP will have a data-structure at every vertex $v \in \mathcal{B}'$. However,
 451 to achieve polylogarithmic time per trial we can only update a polylogarithmic number of these
 452 data-structures per trial. This necessitates the use of *contractions* of our trees $\{\mathcal{Z}_t \mid t \in [T] \setminus \{1\}\}$
 453 which are defined as follows. A *contraction* of a full binary tree \mathcal{Q} is another full binary tree \mathcal{J} which
 454 satisfies the following:

- 455 • The vertices of \mathcal{J} are a subset of those of \mathcal{Q} .
- 456 • $r(\mathcal{J}) = r(\mathcal{Q})$
- 457 • Given an internal vertex $u \in \mathcal{J}^\dagger$ we have $\triangleleft_{\mathcal{J}}(u) \in \downarrow_{\mathcal{Q}}(\triangleleft_{\mathcal{Q}}(u))$ and $\triangleright_{\mathcal{J}}(u) \in \downarrow_{\mathcal{Q}}(\triangleright_{\mathcal{Q}}(u))$
- 458 • Any leaf of \mathcal{J} is a leaf of \mathcal{Q} .

459 Note that any contraction of \mathcal{Z}_t is also a contraction of \mathcal{Z}_{t+1} and hence, by induction, a contraction
 460 of $\mathcal{Z}_{t'}$ for all $t' \geq t$. Given a trial t and a contraction \mathcal{J} of \mathcal{Z}_{t-1} we now define the operation
 461 **INSERT $_t(\mathcal{J})$** which acts on \mathcal{J} by the following algorithm:

- 462 1. Let \hat{u} be the unique vertex in $\mathcal{J} \setminus r(\mathcal{J})$ such that u_t lies in the maximal spanning tree of \mathcal{Z}_t
 463 with $\uparrow_{\mathcal{J}}(\hat{u})$ and \hat{u} as leaves.
- 464 2. Let $u^* := \Gamma_{\mathcal{Z}_t}(u_t, \hat{u})$.
- 465 3. Add the vertices u^* and u_t to the tree \mathcal{J} .
- 466 4. Let $u' := \uparrow_{\mathcal{J}}(\hat{u})$.
- 467 5. If $\hat{u} = \triangleleft_{\mathcal{J}}(u')$ then set $\triangleleft_{\mathcal{J}}(u') \leftarrow u^*$. Else set $\triangleright_{\mathcal{J}}(u') \leftarrow u^*$.
- 468 6. If $\hat{u} \in \downarrow_{\mathcal{Z}_t}(\triangleleft_{\mathcal{Z}_t}(u^*))$ then set $\triangleleft_{\mathcal{J}}(u^*) \leftarrow \hat{u}$ and $\triangleright_{\mathcal{J}}(u^*) \leftarrow u_t$. Else set $\triangleright_{\mathcal{J}}(u^*) \leftarrow \hat{u}$ and
 469 $\triangleleft_{\mathcal{J}}(u^*) \leftarrow u_t$

470 Later in this paper we will show how this operation can be done in polylogarithmic time. Note that
 471 after the operation we have that \mathcal{J} is a contraction of \mathcal{Z}_t and u_t has been added to it's leaves. From
 472 now on when we use the term *contraction* we mean any contraction of \mathcal{Z}_T .

473 D.3 Contraction-Based Bayesian Networks

474 Here we shall define a bayesian network over any contraction \mathcal{J} and show how it can be utilised to
 475 compute certain quantities required by CANPROP. This bayesian network takes a parameter $\epsilon \in [0, 1]$.
 476 First define the quantity $\phi_0(\epsilon) := 0$ and for all $j \in \mathbb{N} \cup \{0\}$ inductively define:

$$\phi_{j+1}(\epsilon) := (1 - \epsilon)\phi_j(\epsilon) + \epsilon(1 - \phi_j(\epsilon))$$

477 The algorithm must compute these quantities for various values of ϵ . However, for all $t \in [T]$ we
 478 have that $\phi_t(\epsilon)$ doesn't have to be computed until trial t so computing these quantities is constant
 479 time per trial (for each value of ϵ). Given a contraction \mathcal{J} , a value $\epsilon \in [0, 1]$, a vertex $u \in \mathcal{J} \setminus r(\mathcal{J})$
 480 and indices $i, i' \in \{0, 1\}$ define:

$$\tilde{\tau}_{i,i'}(\mathcal{J}, u, \epsilon) := \llbracket i \neq i' \rrbracket \phi_{(d(u)-d(\uparrow_{\mathcal{J}}(u)))}(\epsilon) + \llbracket i = i' \rrbracket (1 - \phi_{(d(u)-d(\uparrow_{\mathcal{J}}(u)))}(\epsilon))$$

481 which defines the transition matrix from $\uparrow_{\mathcal{J}}(u)$ to u in a bayesian network. We shall now show
 482 how belief propagation over such bayesian networks can be used to compute the quantities we need
 483 in CANPROP. Suppose we have a contraction \mathcal{J} , a value $\epsilon \in [0, 1]$ and a function $\lambda : \mathcal{J}^* \rightarrow \mathbb{R}_+$.
 484 This function λ induces a function $\lambda' : \mathcal{X} \rightarrow \mathbb{R}_+$ defined as follows. Given $x \in \mathcal{X}$, if there exists a
 485 leaf $u \in \mathcal{J}^*$ with $\gamma(u) = x$ then $\lambda'(x) = \lambda(u)$. Otherwise $\lambda'(x) = 1$. We then define a weighting
 486 $\tilde{w}(\lambda, \epsilon, \cdot) : 2^{\mathcal{X}} \rightarrow \mathbb{R}_+$ such that for all $\mathcal{S} \in 2^{\mathcal{X}}$ we have:

$$\tilde{w}(\lambda, \epsilon, \mathcal{S}) := \left(\prod_{x \in \mathcal{S}} \lambda'(x) \right) \left(\prod_{x \in \mathcal{X}'} (\sigma(x, \mathcal{S})\epsilon + (1 - \sigma(x, \mathcal{S}))(1 - \epsilon)) \right)$$

487 For the CANPROP algorithm we will need to compute

$$\sum_{\mathcal{S} \in 2^{\mathcal{X}}} \llbracket \gamma(\hat{u}) \in \mathcal{S} \rrbracket \tilde{w}(\lambda, \epsilon, \mathcal{S}) \quad (6)$$

488 for some leaf $\hat{u} \in \mathcal{J}^*$. We shall now show how we can compute this quantity via belief propagation
 489 on the bayesian network. In particular we shall construct a quantity $\tilde{\Lambda}(\mathcal{J}, \lambda, \epsilon, u)$ equal to the quantity
 490 in Equation (6). To do this first define the function $\lambda^* : \mathcal{J} \rightarrow \mathbb{R}_+$ so that for all $u \in \mathcal{J}^*$ we have
 491 $\lambda^*(u) = \lambda(u)$ and for all $u \in \mathcal{J}^\dagger$ we have $\lambda^*(u) = 1$. For all vertices $u \in \mathcal{J}$ and all indices
 492 $i \in \{0, 1\}$ define:

$$\tilde{\kappa}_i(\lambda, u) := \llbracket i = 0 \rrbracket + \llbracket i = 1 \rrbracket \lambda^*(u)$$

493 For all $\hat{u} \in \mathcal{J}$ define:

$$\mathcal{F}(\mathcal{J}, \hat{u}) := \{f \in \{0, 1\}^{\mathcal{J}} \mid f(\hat{u}) = 1\}$$

494 and then define:

$$\tilde{\Lambda}(\mathcal{J}, \lambda, \epsilon, \hat{u}) := \sum_{f \in \mathcal{F}(\mathcal{J}, \hat{u})} \prod_{u \in \mathcal{J} \setminus r(\mathcal{J})} \tilde{\tau}_{f(\uparrow_{\mathcal{J}}(u)), f(u)}(\mathcal{J}, u, \epsilon) \tilde{\kappa}_{f(u)}(\lambda, u)$$

495 The equality of this quantity and that given in Equation (6) is given by the following theorem.

496 **Theorem D.1.** *Given a contraction \mathcal{J} , a function $\lambda : \mathcal{J}^* \rightarrow \mathbb{R}_+$, some $\epsilon \in [0, 1]$ and some leaf*
 497 *$\hat{u} \in \mathcal{J}^*$ we have:*

$$\tilde{\Lambda}(\mathcal{J}, \lambda, \epsilon, \hat{u}) = \sum_{\mathcal{S} \in 2^{\mathcal{X}}} \llbracket \gamma(\hat{u}) \in \mathcal{S} \rrbracket \tilde{w}(\lambda, \epsilon, \mathcal{S})$$

498 Note that $\tilde{\Lambda}(\mathcal{J}, \lambda, \epsilon, \hat{u})$ is of the exact form to be solved via belief propagation over \mathcal{J} . However,
 499 belief propagation is still too slow (taking $\Theta(|\mathcal{J}|)$ time) - we will remedy this later.

500 D.4 Cancellation Propagation with Contractions

501 We now describe how to implement CANPROP with contractions. For each $v \in \mathcal{B}'$ we maintain a
 502 contraction $\mathcal{A}(v)$ and a function $\zeta(v, \cdot) : \mathcal{A}(v)^* \rightarrow \mathbb{R}_+$. We initialise with $\mathcal{A}(v)$ identical to \mathcal{Z}_2 and
 503 $\zeta(v, u) = 1$ for both leaves $u \in \mathcal{Z}_2^*$. Via induction over t we will have that at the start of each trial t
 504 we have, for all sets $\mathcal{S} \in 2^{\mathcal{X}}$, that:

$$w_t(v, \mathcal{S}) = \frac{1}{4 \log(T)} \sum_{i \in [\log(T)]} \tilde{w}(\zeta(v, \cdot), 2^i/T, \mathcal{S}) \quad (7)$$

505 On trial t we do as follows. First we update \mathcal{Z}_{t-1} to \mathcal{Z}_t using the algorithm GROW_t . We will perform
506 the necessary modifications to our contractions as we sample the path \mathcal{P}_t . In particular we first
507 set $v_{t,0} \leftarrow r(\mathcal{B})$ and then for each $j \in [\log(K) - 1] \cup \{0\}$ in turn we do as follows. For each
508 $v \in \{\triangleleft(v_{t,j}), \triangleright(v_{t,j})\}$ run $\text{INSERT}_t(\mathcal{A}(v))$ and set $\zeta(v, u_t) \leftarrow 1$. Since $\zeta(v, u_t) = 1$ Equation (7)
509 still holds and hence, by Theorem D.1, we have:

$$\theta_t(v) = \frac{1}{4 \log(T)} \sum_{i \in [\log(T)]} \tilde{\Lambda}(\mathcal{A}(v), \zeta(v, \cdot), 2^i/T, u_t)$$

510 where $\zeta(v, \cdot)$ is the function that maps each $u \in \mathcal{A}(v)$ to $\zeta(v, u)$. After $\theta_t(v)$ has been computed for
511 both $v \in \{\triangleleft(v_{t,j}), \triangleright(v_{t,j})\}$ we can now sample $v_{t,j+1}$.

512 Once we have selected the action a_t we then update the functions $\{\zeta(v, \cdot) \mid \uparrow_{\mathcal{B}}(v) \in \mathcal{P}_t\}$ by setting
513 $\zeta(v, u_t) \leftarrow \beta_t(v)$ for all $v \in \mathcal{B}'$ with $\uparrow_{\mathcal{B}}(v) \in \mathcal{P}_t$. It is clear now that Equation (7) holds inductively.

514 D.5 Notational Relationship to the Main Body

515 We now point out how the notation in this section relates to that of the main body. In particular we
516 have, for all $v \in \mathcal{B}'$, all $u \in \mathcal{A}(v)$, all $j \in [\log(T)]$ and all $i, i' \in \{0, 1\}$, that:

- 517 • $\tau_{i,i'}(\mathcal{A}(v), u, j) = \tilde{\tau}_{i,i'}(\mathcal{A}(v), u, 2^j/T)$
- 518 • $\kappa_i(\mathcal{A}(v), u) = \tilde{\kappa}_i(\zeta(v, \cdot), u)$
- 519 • $\Lambda(\mathcal{A}(v), u, j) = \tilde{\Lambda}(\mathcal{A}(v), \zeta(v, \cdot), 2^j/T, u)$

520 E Utilising Ternary Search Trees

521 There are now only two things left to do in order to achieve polylogarithmic time per trial - to make
522 an efficient online implementation of the $\text{INSERT}_t(\cdot)$ operation and an efficient online algorithm to
523 perform belief propagation over our contractions. In order to do this we will utilise the methodology
524 of [19] which we now describe. However, we do not give the full details of the rebalancing technique
525 and refer the reader to [19] for these details.

526 E.1 Ternary Search Trees

527 In this section we will consider a full binary tree \mathcal{J} . A (full) ternary tree \mathcal{D} is a rooted tree in which
528 each internal vertex $s \in \mathcal{D}^\dagger$ has three children denoted by $\triangleleft(s)$, $\nabla(s)$, $\triangleright(s)$ and called the left, centre,
529 and right children respectively. We now define what it means for a ternary tree \mathcal{D} to be a ternary
530 search tree (TST) of \mathcal{J} . Firstly, the vertex set of \mathcal{D} is partitioned into two sets \mathcal{D}° and \mathcal{D}^\bullet . Every
531 vertex $s \in \mathcal{D}$ is associated with a vertex $\mu(s) \in \mathcal{J}$ and every $s \in \mathcal{D}^\bullet$ is also associated with a vertex
532 $\mu'(s) \in \downarrow_{\mathcal{J}}(\mu(s))^\dagger$. The root $r(\mathcal{D})$ of \mathcal{D} is contained in \mathcal{D}° and $\mu(r(\mathcal{D})) := r(\mathcal{J})$. Each internal
533 vertex $s \in \mathcal{D}^\dagger$ is associated with a vertex $\xi(s) \in \mathcal{J}$. If $s \in \mathcal{D}^\circ$ then $\xi(s) \in \downarrow(\mu(s))^\dagger$ and if $s \in \mathcal{D}^\bullet$
534 then $\xi(s)$ lies on the path (in \mathcal{J}) from $\mu(s)$ to $\uparrow(\mu'(s))$. For all $s \in \mathcal{D}^\dagger$ we have:

- 535 • $\nabla(s) \in \mathcal{D}^\bullet$, $\mu(\nabla(s)) := \mu(s)$ and $\mu'(\nabla(s)) := \xi(s)$.
- 536 • $\triangleleft(s)$ satisfies:
 - 537 – If $s \in \mathcal{D}^\circ$ then $\triangleleft(s) \in \mathcal{D}^\circ$ and $\mu(\triangleleft(s)) := \triangleleft(\xi(s))$.
 - 538 – If $s \in \mathcal{D}^\bullet$ and $\mu'(s) \in \downarrow(\triangleright(\xi(s)))$ then $\triangleleft(s) \in \mathcal{D}^\circ$ and $\mu(\triangleleft(s)) := \triangleleft(\xi(s))$
 - 539 – Else $\triangleleft(s) \in \mathcal{D}^\bullet$, $\mu(\triangleleft(s)) := \triangleleft(\xi(s))$ and $\mu'(\triangleleft(s)) := \mu'(s)$
- 540 • $\triangleright(s)$ satisfies:
 - 541 – If $s \in \mathcal{D}^\circ$ then $\triangleright(s) \in \mathcal{D}^\circ$ and $\mu(\triangleright(s)) := \triangleright(\xi(s))$.
 - 542 – If $s \in \mathcal{D}^\bullet$ and $\mu'(s) \in \downarrow(\triangleleft(\xi(s)))$ then $\triangleright(s) \in \mathcal{D}^\circ$ and $\mu(\triangleright(s)) := \triangleright(\xi(s))$
 - 543 – Else $\triangleright(s) \in \mathcal{D}^\bullet$, $\mu(\triangleright(s)) := \triangleright(\xi(s))$ and $\mu'(\triangleright(s)) := \mu'(s)$

544 Finally, for each leaf $s \in \mathcal{D}^*$ we have:

- 545 • If $s \in \mathcal{D}^\circ$ then $\mu(s)$ is a leaf of \mathcal{J} .

546 • If $s \in \mathcal{D}^\bullet$ then there exists $u \in \mathcal{J}^\dagger$ such that $\mu(s) = \mu'(s) = u$.

547 Intuitively each vertex $s \in \mathcal{D}$ is associated with a subtree $\hat{\mathcal{J}}(s)$ of \mathcal{J} . If $s \in \mathcal{D}^\circ$ then $\hat{\mathcal{J}}(s) := \Downarrow(\mu(s))$
 548 and if $s \in \mathcal{D}^\bullet$ then $\hat{\mathcal{J}}(s)$ is the subtree of descendants of $\mu(s)$ which are not proper descendants
 549 of $\mu'(s)$. For every $s \in \mathcal{D}$ such that $\hat{\mathcal{J}}(s)$ contains only a single vertex, we have that s is a leaf of
 550 \mathcal{D} . Otherwise s is an internal vertex of \mathcal{D} and its children are as follows. We say that $\hat{\mathcal{J}}(s)$ is *split*
 551 at the vertex $\xi(s) \in \hat{\mathcal{J}}(s)^\dagger$. If $s \in \mathcal{D}^\bullet$ we require that $\xi(s)$ is on the path in \mathcal{J} from $\mu(s)$ to $\mu'(s)$.
 552 The action of splitting $\hat{\mathcal{J}}(s)$ at $\xi(s)$ partitions $\hat{\mathcal{J}}(s)$ into the subtrees $\hat{\mathcal{J}}(\triangleleft(s))$, $\hat{\mathcal{J}}(\nabla(s))$ and $\hat{\mathcal{J}}(\triangleright(s))$
 553 defined as follows:

- 554 • $\hat{\mathcal{J}}(\triangleleft(s)) := \Downarrow(\triangleleft(\xi(s))) \cap \hat{\mathcal{J}}(s)$
 555 • $\hat{\mathcal{J}}(\triangleright(s)) := \Downarrow(\triangleright(\xi(s))) \cap \hat{\mathcal{J}}(s)$
 556 • $\hat{\mathcal{J}}(\nabla(s)) := \hat{\mathcal{J}}(s) \setminus (\hat{\mathcal{J}}(\triangleleft(s)) \cup \hat{\mathcal{J}}(\triangleright(s)))$

557 Utilising the methodology of [19] we will maintain TSTs of \mathcal{Z}_t (at each trial t) and the trees in
 558 $\{\mathcal{A}(v) \mid v \in \mathcal{B}'\}$, each with height $\mathcal{O}(\ln(T))$. Note that these trees are dynamic, in that vertices are
 559 inserted into them over time. [19] shows how, after such an insertion, the corresponding TST can be
 560 *rebalanced* so that its height is still in $\mathcal{O}(\ln(T))$. This rebalancing is performed via a sequence of
 561 $\mathcal{O}(\ln(T))$ *tree rotations*, which generalise the concept of tree rotations in binary search trees.

562 E.2 Searching

563 In this section we show how we can use our TSTs to implement the operation $\text{INSERT}_t(\mathcal{J})$ on any
 564 trial t and contraction \mathcal{J} of \mathcal{Z}_{t-1} . To do this we need to perform the following two search operations:

- 565 1. Find the unique vertex $\hat{u} \in \mathcal{J} \setminus r(\mathcal{J})$ such that u_t lies in the maximal spanning tree of \mathcal{Z}_t
 566 with $\uparrow_{\mathcal{J}}(\hat{u})$ and \hat{u} as leaves.
 567 2. Find $u^* := \Gamma_{\mathcal{Z}_t}(u_t, \hat{u})$

568 To perform these tasks in polylogarithmic time we will utilise TSTs \mathcal{E} and \mathcal{D} of \mathcal{Z}_t and \mathcal{J} respectively.
 569 Both the searching tasks utilise a function $\nu : \mathcal{Z}_t^2 \rightarrow \{\blacktriangle, \blacktriangleleft, \blacktriangleright\}$ defined, for all $u, u' \in \mathcal{Z}_t$ as follows.
 570 If $u' \in \Downarrow_{\mathcal{Z}_t}(\triangleleft(u))$ or $u' \in \Downarrow_{\mathcal{Z}_t}(\triangleright(u))$ then $\nu(u, u') := \blacktriangleleft$ or $\nu(u, u') := \blacktriangleright$ respectively. Otherwise
 571 $\nu(u, u') := \blacktriangle$. This can be computed as follows. If $u = u'$ then $\nu(u, u') := \blacktriangle$. Otherwise let \tilde{s} and
 572 \tilde{s}' be the unique leaves of \mathcal{E} such that $\mu(\tilde{s}) = u$ and $\mu(\tilde{s}') = u'$. Let $s^* := \Gamma_{\mathcal{E}}(\tilde{s}, \tilde{s}')$ and let \hat{s} and
 573 \hat{s}' be the children of s^* which are ancestors of \tilde{s} and \tilde{s}' respectively. If $\hat{s} \neq \nabla(s^*)$ then we have
 574 $\nu(u, u') = \blacktriangle$. If $\xi(s^*) = u$ then we have $\nu(u, u') = \blacktriangleleft$ or $\nu(u, u') = \blacktriangleright$ if $\hat{s}' = \triangleleft(s^*)$ or $\hat{s}' = \triangleright(s^*)$
 575 respectively. If $\hat{s} = \nabla(s^*)$ and $\xi(s^*) \neq u$ then we perform the following process. Start with s
 576 equal to \hat{s} . At any point in the process we do as follows. If $s \in \mathcal{E}^\circ$ then the process terminates
 577 with $\nu(u, u') := \blacktriangle$. If $s \in \mathcal{E}^\bullet$ and $u = \xi(s)$ then the process terminates with $\nu(u, u') = \blacktriangleleft$ or
 578 $\nu(u, u') = \blacktriangleright$ if $\triangleleft(s) \in \mathcal{E}^\bullet$ or $\triangleright(s) \in \mathcal{E}^\bullet$ respectively. If $s \in \mathcal{E}^\bullet$ and $u \neq \xi(s)$ then we reset s as equal
 579 to the child of s which is an ancestor of \tilde{s} and continue the process.

580 The vertex \hat{u} can be found as follows. We construct a root-to-leaf path in \mathcal{D} such that, given a vertex
 581 s in the path, the next vertex in the path is $\triangleleft(s)$, $\triangleright(s)$ or $\nabla(s)$ if $\nu(\xi(s), u_t)$ is equal to \blacktriangleleft , \blacktriangleright or \blacktriangle
 582 respectively. Given that s' is the leaf of \mathcal{D} that is in this path we have $\hat{u} = \mu(s')$.

583 The vertex u^* can then be found as follows. We construct a root-to-leaf path in \mathcal{E} such that, given a
 584 vertex s in the path, the next vertex in the path is found as follows. If $\nu(\xi(s), u_t) = \nu(\xi(s), \hat{u})$ then
 585 given $\nu(\xi(s), u_t)$ is equal to \blacktriangleleft , \blacktriangleright or \blacktriangle , the next vertex is equal to $\triangleleft(s)$, $\triangleright(s)$ or $\nabla(s)$ respectively.
 586 Otherwise, the next vertex is $\nabla(s)$. Given that s' is the leaf of \mathcal{E} that is in this path we have $u^* = \mu(s')$.

587 The fact that these algorithms find the correct vertices is given in the following theorem:

588 **Theorem E.1.** *The above algorithms are correct.*

589 E.3 Belief Propagation

590 Here we utilise the methodology of [6] in order to efficiently compute the function $\tilde{\Lambda}$ that appears
 591 in the CANPROP implementation. i.e. given a contraction \mathcal{J} , a function $\lambda : \mathcal{J}^* \rightarrow \mathbb{R}_+$, some

592 $\epsilon \in [0, 1]$ and some leaf $\hat{u} \in \mathcal{J}^*$ we need to compute $\tilde{\Lambda}(\mathcal{J}, \lambda, \epsilon, \hat{u})$. For brevity let us define, for all
 593 $i, i' \in \{0, 1\}$, and all vertices $u \in \mathcal{J} \setminus \{r(\mathcal{J})\}$, the quantities:

$$\hat{\tau}_{i,i'}(u) := \tilde{\tau}_{i,i'}(\mathcal{J}, u, \epsilon) \quad ; \quad \hat{\kappa}_i(u) := \tilde{\kappa}_i(\lambda, u)$$

594 For simplicity of presentation we will utilise a tree \mathcal{J}' which is defined as identical to \mathcal{J} except with
 595 a single vertex added as the parent of $r(\mathcal{J})$. For all $i, i' \in \{0, 1\}$ we define $\hat{\kappa}_i(r(\mathcal{J}')) := 1$ and
 596 $\hat{\tau}_{i,i'}(r(\mathcal{J})) = \llbracket i = i' \rrbracket$. For all $u \in \mathcal{J}$ we will define $\uparrow(u) := \uparrow_{\mathcal{J}'}(u)$

597 We will utilise a TST \mathcal{D} of \mathcal{J} by maintaining *potentials* on the vertices of \mathcal{D} defined as follows. First,
 598 for any vertex $s \in \mathcal{D}$ define the subtree $\tilde{\mathcal{J}}(s)$ of \mathcal{J} to be equal to $\downarrow_{\mathcal{J}}(\mu(s))$ if $s \in \mathcal{D}^\circ$ and equal
 599 to the maximal subtree with $\mu(s)$ and $\mu'(s)$ as leaves if $s \in \mathcal{D}^\bullet$. For all $s \in \mathcal{D}^\circ$ and $i \in \{0, 1\}$ we
 600 define:

$$\Psi_i(s) := \sum_{f \in \{0,1\}^{\tilde{\mathcal{J}}(s) \cup \{\uparrow(\mu(s))\}}} \llbracket f(\uparrow(\mu(s))) = i \rrbracket \prod_{u \in \tilde{\mathcal{J}}(s)} \hat{\tau}_{f(\uparrow(u)), f(u)}(u) \hat{\kappa}_{f(u)}(u)$$

601 and for all $s \in \mathcal{D}^\bullet$ and $i, i' \in \{0, 1\}$ we define:

$$\Omega_{i,i'}(s) := \sum_{f \in \{0,1\}^{\tilde{\mathcal{J}}(s) \cup \{\uparrow(\mu(s))\}}} \llbracket f(\uparrow(\mu(s))) = i \rrbracket \llbracket f(\mu'(s)) = i' \rrbracket \prod_{u \in \tilde{\mathcal{J}}(s)} \hat{\tau}_{f(\uparrow(u)), f(u)}(u) \hat{\kappa}_{f(u)}(u)$$

602 We have the following recurrence relations for these potentials. Suppose we have an internal vertex
 603 $s \in \mathcal{D}^\dagger$ and $i, i' \in \{0, 1\}$. If $s \in \mathcal{D}^\circ$ we have:

$$\Psi_i(s) = \sum_{i'' \in \{0,1\}} \Omega_{i,i''}(\nabla(s)) \Psi_{i''}(\triangleleft(s)) \Psi_{i''}(\triangleright(s))$$

604 If, instead, $s \in \mathcal{D}^\bullet$ then, by letting $s' := \triangleleft(s)$, $s'' := \triangleright(s)$ if $\triangleleft(s) \in \mathcal{D}^\bullet$ and $s' := \triangleright(s)$, $s'' := \triangleleft(s)$
 605 otherwise, we have:

$$\Omega_{i,i'}(s) = \sum_{i'' \in \{0,1\}} \Omega_{i,i''}(\nabla(s)) \Omega_{i'',i'}(s') \Psi_{i''}(s'')$$

606 If, on a trial t , we perform the operation $\text{INSERT}_t(\mathcal{J})$ or change the value of $\lambda(u_t)$ these recurrence
 607 relations can be used to update the potentials (in conjunction with the tree rotations) in logarithmic
 608 time.

609 Now that we have defined our potentials we will show how to use them to compute $\tilde{\Lambda}(\mathcal{J}, \lambda, \epsilon, \hat{u})$
 610 in logarithmic time. To do this we recursively define the following quantities for $i \in \{0, 1\}$. Let
 611 $\omega_i(r(\mathcal{D})) := 1$. Given an internal vertex $s \in \mathcal{D}^\circ$ we define:

$$\omega_i(\nabla(s)) := \omega_i(s) \quad ; \quad \omega'_i(\nabla(s)) := \Psi_i(\triangleleft(s)) \Psi_i(\triangleright(s))$$

612

$$\omega_i(\triangleleft(s)) := \Psi_i(\triangleright(s)) \sum_{i' \in \{0,1\}} \omega_{i'}(s) \Omega_{i',i}(s) \quad ; \quad \omega_i(\triangleright(s)) := \Psi_i(\triangleleft(s)) \sum_{i' \in \{0,1\}} \omega_{i'}(s) \Omega_{i',i}(s)$$

613 Given an internal vertex $s \in \mathcal{D}^\bullet$ define $s' := \triangleleft(s)$, $s'' := \triangleright(s)$ if $\triangleleft(s) \in \mathcal{D}^\bullet$ and $s' := \triangleright(s)$,
 614 $s'' := \triangleleft(s)$ otherwise. Then:

$$\omega_i(\nabla(s)) := \omega_i(s) \quad ; \quad \omega'_i(\nabla(s)) := \Psi_i(s'') \sum_{i' \in \{0,1\}} \Omega_{i,i'}(s') \omega'_{i'}(s)$$

615

$$\omega_i(s') := \sum_{i' \in \{0,1\}} \omega_{i'}(s) \Omega_{i',i}(s) \Psi_i(s'') \quad ; \quad \omega'_i(s') := \omega'_i(s)$$

616

$$\omega_i(s'') := \sum_{i', i'' \in \{0,1\}} \omega_{i'}(s) \Omega_{i',i}(\nabla(s)) \omega'_{i''}(s) \Omega_{i,i''}(s')$$

617 For $s \in \mathcal{D}^\circ$, $\omega'_i(s)$ is not required and hence is arbitrary. We inductively compute the values
 618 $\{\omega_i(s), \omega'_i(s) \mid i \in \{0, 1\}\}$ for all s in the path from $r(\mathcal{D})$ to the unique leaf $\hat{s} \in \mathcal{D}^*$ in which
 619 $\mu(\hat{s}) = \hat{u}$. We then have $\tilde{\Lambda}(\mathcal{J}, \lambda, \epsilon, \hat{u}) = \omega_1(\hat{s})$.

620 Since this is known methodology we do not include a proof in this paper and direct the reader to [6].

621 **F Proofs**

622 **F.1 Theorem 3.1**

623 For brevity we write α instead of $\alpha(y, \mu)$. Choose some $\delta > 0$. Let $\mathcal{E} := \mathcal{C} \setminus \mathcal{M}(y, \delta)$. Let \mathcal{X} be a
 624 set of m contexts drawn i.i.d. at random from μ . Now consider some x drawn from μ and let \hat{x} be a
 625 c -nearest neighbour of x in \mathcal{X} .

626 Suppose that $x \in \mathcal{E}$. Let \mathcal{A} be the ball of radius δ/c centred at x . We have that:

$$\mu(\mathcal{A}) \geq \epsilon \lambda \left(\frac{\delta}{c} \right)^d$$

627 where λ is a constant dependent on d . This means that for any x' drawn from μ we have that:

$$\mathbb{P}[x' \notin \mathcal{A}] \leq 1 - \lambda \epsilon \left(\frac{\delta}{c} \right)^d \leq \exp \left(-\lambda \epsilon \left(\frac{\delta}{c} \right)^d \right)$$

628 Suppose that:

$$m \geq \frac{-\ln(\alpha \delta)}{\lambda \epsilon} \left(\frac{c}{\delta} \right)^d$$

629 Note that if there exists $x'' \in \mathcal{X}$ with $x'' \in \mathcal{A}$ then $\Delta(x, \hat{x}) \leq \delta$ so that $y(x) = y(\hat{x})$. The above
 630 equations then give us:

$$\mathbb{P}[y(x) \neq y(\hat{x}) | x \in \mathcal{E}] \leq \exp \left(-m \lambda \epsilon \left(\frac{\delta}{c} \right)^d \right) \leq \alpha \delta$$

631 We then have that:

$$\mathbb{P}[y(x) \neq y(\hat{x})] \leq \alpha \delta + \mu(\mathcal{M}(\delta)) \in \mathcal{O}(\alpha \delta)$$

632 Since:

$$\delta \in \tilde{\mathcal{O}}(c(\epsilon m)^{-1/d})$$

633 we now have:

$$\mathbb{P}[y(x) \neq y(\hat{x})] \in \tilde{\mathcal{O}} \left(c \alpha (\epsilon m)^{-1/d} \right)$$

634 **F.2 Theorem 3.2**

635 This theorem is proved in appendices C to E and the theorems therein.

636 **F.3 Theorem 3.3**

637 Choose a set $\mathcal{S} \subseteq \mathcal{C}$ in which for all $t \in [T]$ there exists $x \in \mathcal{S}$ with $\Delta(x, x_t) < \gamma(x, y)/3c$. For all
 638 trials t let \mathcal{S}_t be the set of all contexts $x \in \mathcal{S}$ in which there exists $s \in [t]$ with $\Delta(x, x_s) < \gamma(x, y)/3c$.

639 Now consider a trial t in which $y(x_t) \neq y(n(x_t))$ and choose $x \in \mathcal{S}$ with $\Delta(x, x_t) < \gamma(x, y)/3c$.

640 Assume, for contradiction, that $x \in \mathcal{S}_{t-1}$. Then there exists $s \in [t-1]$ with $\Delta(x, x_s) < \gamma(x, y)/3c$
 641 so that by the triangle inequality we have:

$$\Delta(x_t, x_s) \leq \Delta(x, x_s) + \Delta(x, x_t) < 2\gamma(x, y)/3c$$

642 which implies that $\Delta(x_t, n(x_t)) < 2\gamma(x, y)/3$. By the triangle inequality we then have that:

$$\Delta(x, n(x_t)) \leq \Delta(x_t, n(x_t)) + \Delta(x, x_t) < 2\gamma(x, y)/3 + \gamma(x, y)/3c \leq 3\gamma(x, y)/3 = \gamma(x, y)$$

643 Since $\Delta(x, x_t) < \gamma(x, y)$ we have $y(x) = y(x_t)$ and hence that $y(x) \neq y(n(x_t))$. But this
 644 contradicts the fact that $\Delta(x, n(x_t)) < \gamma(x, y)$.

645 We have hence shown that $x \notin \mathcal{S}_{t-1}$. Since $x \in \mathcal{S}_t$ we then have that $|\mathcal{S}_t| \geq |\mathcal{S}_{t-1}|$. This implies
 646 that:

$$\Phi(y) = \sum_{t \in [T]} \mathbb{I}[y(x_t) \neq y(n(x_t))] \leq |\mathcal{S}_T| \leq |\mathcal{S}|$$

647 as required.

648 **F.4 Theorem 3.4**

649 By linearity of expectation we have:

$$\mathbb{E}[\Phi(y)] \leq 1 + \sum_{t \in [T]} g_t(\mu, y, c)$$

650 and from Theorem 3.1 we have:

$$g_t(\mu, y, c) \in \mathcal{O}\left(c\alpha(y, \mu)\epsilon t^{-1/d}\right)$$

651 so that:

$$\mathbb{E}[\Phi(y)] \in \mathcal{O}\left(c\alpha(y, \mu)\epsilon^{-1/d}T^{(d-1)/d}\ln(T)\right)$$

652 By setting:

$$\rho := T^{(d-1)/(2d)}c^{1/2}$$

653 we then have:

$$\rho + \frac{\mathbb{E}[\Phi(y)]}{\rho} \in \mathcal{O}\left(c^{1/2}(1 + \alpha(y, \mu)\epsilon^{-1/d})T^{(d-1)/(2d)}\ln(T)\right)$$

654 so that by Theorem 3.2 we have:

$$\mathbb{E}[R(y)] \in \tilde{\mathcal{O}}\left(c^{1/2}(1 + \alpha(y, \mu)\epsilon^{-1/d})K^{1/2}T^{(2d-1)/(2d)}\right)$$

655 **F.5 Theorem C.1**

656 For every trial $t \in [T]$ define:

$$\Delta_t := - \sum_{v \in \mathcal{B}'} \mathbb{I}[\mathcal{Q}(v) \neq \emptyset] \ln(w_t(v, \mathcal{Q}(v)))$$

657 Choose some arbitrary trial $t \in [T]$. From here until we say otherwise all probabilities and expectations (i.e. whenever we use $\mathbb{P}[\cdot]$ or $\mathbb{E}[\cdot]$) are implicitly conditional on the state of the algorithm at the start of trial t . Note first that we have:

$$\Delta_t - \Delta_{t+1} = \sum_{v \in \mathcal{B}'} \mathbb{I}[\mathcal{Q}(v) \neq \emptyset] \ln\left(\frac{w_{t+1}(v, \mathcal{Q}(v))}{w_t(v, \mathcal{Q}(v))}\right) \quad (8)$$

660 For all $j \in [\log(K)] \cup \{0\}$ let $\gamma_{t,j}$ be the ancestor (in \mathcal{B}) of $y(x_t)$ at depth j . Note that for all $v \in \mathcal{X} \setminus \{\gamma_{t,j} \mid j \in [\log(K)] \cup \{0\}\}$ we have $y(x_t) \notin \downarrow(v)$ so that $x_t \notin \mathcal{Q}(v)$ and hence, directly from the CANPROP algorithm, we have $w_{t+1}(v, \mathcal{Q}(v)) = w_t(v, \mathcal{Q}(v))$. By Equation (8) and the fact that $\mathcal{Q}(v) \neq \emptyset$ for all ancestors v of $y(x_t)$ this implies that:

$$\Delta_t - \Delta_{t+1} = \sum_{j \in [\log(K)]} \ln\left(\frac{w_{t+1}(\gamma_{t,j}, \mathcal{Q}(\gamma_{t,j}))}{w_t(\gamma_{t,j}, \mathcal{Q}(\gamma_{t,j}))}\right) \quad (9)$$

664 For all $j \in [\log(K)]$ define:

$$\lambda_{t,j} := \ln\left(\frac{w_{t+1}(\gamma_{t,j}, \mathcal{Q}(\gamma_{t,j}))}{w_t(\gamma_{t,j}, \mathcal{Q}(\gamma_{t,j}))}\right)$$

665 and:

$$\epsilon_{t,j} := \mathbb{E}[\ln(\psi_{t,j}) \mid \gamma_{t,j} \in \mathcal{P}_t]$$

666 Now choose some arbitrary $j \in [\log(K)]$. If $\gamma_{t,(j-1)} \in \mathcal{P}_t$ then $\gamma_{t,(j-1)} = v_{t,(j-1)}$ so $\uparrow(\gamma_{t,j}) = v_{t,(j-1)}$ and hence, since $x_t \in \mathcal{Q}(\gamma_{t,j})$, we have $\lambda_{t,j} = \ln(\beta_t(\gamma_{t,j}))$. By definition of $\beta_t(\gamma_{t,j})$ this means that:

$$\mathbb{E}[\lambda_{t,j} \mid \gamma_{t,j} \in \mathcal{P}_t, \gamma_{t,(j-1)} \in \mathcal{P}_t] = \epsilon_{t,j} - \mathbb{E}[\ln(\psi_{t,(j-1)}) \mid \gamma_{t,j} \in \mathcal{P}_t, \gamma_{t,(j-1)} \in \mathcal{P}_t]$$

669 and that:

$$\mathbb{E}[\lambda_{t,j} \mid \gamma_{t,j} \notin \mathcal{P}_t, \gamma_{t,(j-1)} \in \mathcal{P}_t] = -\mathbb{E}[\ln(\psi_{t,(j-1)}) \mid \gamma_{t,j} \notin \mathcal{P}_t, \gamma_{t,(j-1)} \in \mathcal{P}_t]$$

670 Multiplying these two equations by $\mathbb{P}[\gamma_{t,j} \in \mathcal{P}_t \mid \gamma_{t,(j-1)} \in \mathcal{P}_t]$ and $\mathbb{P}[\gamma_{t,j} \notin \mathcal{P}_t \mid \gamma_{t,(j-1)} \in \mathcal{P}_t]$
671 respectively, and summing them together, then gives us:

$$\mathbb{E}[\lambda_{t,j} \mid \gamma_{t,(j-1)} \in \mathcal{P}_t] = \mathbb{P}[\gamma_{t,j} \in \mathcal{P}_t \mid \gamma_{t,(j-1)} \in \mathcal{P}_t] \epsilon_{t,j} - \mathbb{E}[\ln(\psi_{t,(j-1)}) \mid \gamma_{t,(j-1)} \in \mathcal{P}_t]$$

672 Since $\mathbb{P}[\gamma_{t,j} \in \mathcal{P}_t \mid \gamma_{t,(j-1)} \in \mathcal{P}_t] = \pi_t(\gamma_{t,j})$ we then have:

$$\mathbb{E}[\lambda_{t,j} \mid \gamma_{t,(j-1)} \in \mathcal{P}_t] = \pi_t(\gamma_{t,j}) \epsilon_{t,j} - \epsilon_{t,(j-1)} \quad (10)$$

673 If, on the other hand, $\gamma_{t,(j-1)} \notin \mathcal{P}_t$ then $\uparrow(\gamma_{t,j}) \notin \mathcal{P}_t$ so $\lambda_{t,j} = 0$. This means that:

$$\mathbb{E}[\lambda_{t,j}] = \mathbb{P}[\gamma_{t,(j-1)} \in \mathcal{P}_t] \mathbb{E}[\lambda_{t,j} \mid \gamma_{t,(j-1)} \in \mathcal{P}_t] \quad (11)$$

674 Since the probability that $\gamma_{t,(j-1)} \in \mathcal{P}_t$ is equal to $\prod_{j' \in [j-1]} \pi_t(\gamma_{t,j'})$ we then have, by combining
675 equations (10) and (11), that:

$$\mathbb{E}[\lambda_{t,j}] = \epsilon_{t,j} \prod_{j' \in [j]} \pi_t(\gamma_{t,j'}) - \epsilon_{t,(j-1)} \prod_{j' \in [j-1]} \pi_t(\gamma_{t,j'})$$

676 By substituting into Equation (9) (after taking expectations) we then have that:

$$\begin{aligned} \mathbb{E}[\Delta_t - \Delta_{t+1}] &= -\epsilon_{t,0} + \epsilon_{t,\log(K)} \prod_{j \in [\log(K)]} \pi_t(\gamma_{t,j}) \\ &= -\mathbb{E}[\ln(\psi_{t,0})] + \mathbb{E}[\ln(\psi_{t,\log(K)}) \mid a_t = \gamma_{t,\log(K)}] \prod_{j \in [\log(K)]} \pi_t(\gamma_{t,j}) \end{aligned} \quad (12)$$

677 Note that if $a_t = \gamma_{t,\log(K)}$ then $\gamma_{t,j} = v_{t,j}$ for all $j \in [\log(K)]$. By definition of $\psi_{t,\log(K)}$ and the
678 fact that $\gamma_{t,\log(K)} = y(x_t)$, Equation (12) then gives us:

$$\mathbb{E}[\Delta_t - \Delta_{t+1}] = -\mathbb{E}[\ln(\psi_{t,0})] - \eta \ell_{t,y(x_t)} \quad (13)$$

679 For all $(v, a) \in \mathcal{B} \times [K]$ define:

$$p_{t,a}(v) = \mathbb{P}[a_t = a \mid v \in \mathcal{P}_t]$$

680 noting that this is non-zero only when $a \in \Downarrow(v)$. Suppose we have some $v \in \mathcal{B} \setminus \{r(\mathcal{B})\}$ and some
681 $a \in \Downarrow(v) \cap [K]$. Then, since $\mathbb{P}[a_t = a \mid v \notin \mathcal{P}_t] = 0$, we have:

$$p_{t,a}(\uparrow(v)) = \mathbb{P}[a_t = a \mid \uparrow(v) \in \mathcal{P}_t] = \mathbb{P}[a_t = a \mid v \in \mathcal{P}_t] \mathbb{P}[v \in \mathcal{P}_t \mid \uparrow(v) \in \mathcal{P}_t] = \pi_t(v) p_{t,a}(v)$$

682 Since $p_{t,a}(v) = 0$ whenever $a \notin \Downarrow(v)$, this implies that for all $(v, a) \in \mathcal{B}^\dagger \times [K]$ we have:

$$p_{t,a}(v) = \pi_t(\triangleleft(v)) p_{t,a}(\triangleleft(v)) + \pi_t(\triangleright(v)) p_{t,a}(\triangleright(v)) \quad (14)$$

683 For all $a \in [K]$ define:

$$\hat{\ell}_{t,a} = \frac{\mathbb{I}[a_t = a] \ell_{t,a}}{\mathbb{P}[a_t = a]}$$

684 We now take the inductive hypothesis that for all $j \in [\log(K)] \cup \{0\}$ we have:

$$\psi_{t,j} = \sum_{a \in [K]} p_{t,a}(v_{t,j}) \exp(-\eta \hat{\ell}_{t,a})$$

685 and prove this via reverse induction (i.e. from $j = \log(K)$ to $j = 0$). Note that given $a' := a_t$ we
686 have $\mathbb{P}[a_t = a'] = \prod_{j \in [\log(K)]} \pi_t(v_{t,j})$ and hence:

$$\psi_{t,\log(K)} = \exp(-\eta \hat{\ell}_{t,a_t})$$

687 so the inductive hypothesis holds for $j = \log(K)$. Now suppose that we have some $j' \in [\log(K)]$ and
688 that the inductive hypothesis holds for $j = j'$. We shall now show that it holds also for $j = j' - 1$. Let
689 v' be the child of $v_{t,(j'-1)}$ that is not equal to $v_{t,j'}$. Note that $a_t \notin \Downarrow(v')$ and hence $\exp(-\eta \hat{\ell}_{t,a}) = 1$
690 for all $a \in \Downarrow(v')$ (i.e. whenever $p_{t,a}(v') \neq 0$) which implies:

$$\sum_{a \in [K]} p_{t,a}(v') \exp(-\eta \hat{\ell}_{t,a}) = 1 \quad (15)$$

691 For all $a \in [K]$, Equation (14) gives us:

$$p_{t,a}(v_{t,(j'-1)}) \exp(-\eta \hat{\ell}_{t,a}) = \pi_t(v') p_{t,a}(v') \exp(-\eta \hat{\ell}_{t,a}) + \pi_t(v_{t,j'}) p_{t,a}(v_{t,j'}) \exp(-\eta \hat{\ell}_{t,a})$$

692 Substituting Equation (15) and the inductive hypothesis into this equation (when summed over all
693 $a \in [K]$) then gives us:

$$\sum_{a \in [K]} p_{t,a}(v_{t,(j'-1)}) \exp(-\eta \hat{\ell}_{t,a}) = \pi_t(v') + \pi_t(v_{t,j'}) \psi_{t,j'}$$

694 Since $\pi_t(v') + \pi_t(v_{t,j'}) = 1$ we have, direct from the algorithm, that $\pi_t(v') + \pi_t(v_{t,j'}) \psi_{t,j'} =$
695 $\psi_{t,(j'-1)}$ so the inductive hypothesis holds for $j = j' - 1$. We have hence shown that the inductive
696 hypothesis holds for all $j \in [\log(K)] \cup \{0\}$ and in particular for $j = 0$. Since $p_{t,a}(v_{t,0}) = \mathbb{P}[a_t = a]$
697 we then have:

$$\psi_{t,0} = \sum_{a \in [K]} \mathbb{P}[a_t = a] \exp(-\eta \hat{\ell}_{t,a}) \quad (16)$$

698 Since $\exp(-z) \leq 1 - z + z^2/2$ for all $z \in \mathbb{R}_+$ we have, from Equation (16), that:

$$\psi_{t,0} \leq \sum_{a \in [K]} \mathbb{P}[a_t = a] \left(1 - \eta \hat{\ell}_{t,a} + \frac{\eta^2 \hat{\ell}_{t,a}^2}{2} \right) = 1 - \eta \sum_{a \in [K]} \mathbb{P}[a_t = a] \hat{\ell}_{t,a} + \frac{\eta^2}{2} \sum_{a \in [K]} \mathbb{P}[a_t = a] \hat{\ell}_{t,a}^2$$

699 so since $\ln(1+z) \leq z$ for all $z \in \mathbb{R}$ we have:

$$\ln(\psi_{t,0}) \leq -\eta \sum_{a \in [K]} \mathbb{P}[a_t = a] \hat{\ell}_{t,a} + \frac{\eta^2}{2} \sum_{a \in [K]} \mathbb{P}[a_t = a] \hat{\ell}_{t,a}^2 \quad (17)$$

700 Noting that $\mathbb{P}[a_t = a] \hat{\ell}_{t,a} = \mathbb{I}[a_t = a] \ell_{t,a}$ for all $a \in [K]$, we have:

$$\mathbb{E} \left[\sum_{a \in [K]} \mathbb{P}[a_t = a] \hat{\ell}_{t,a} \right] = \mathbb{E}[\ell_{t,a_t}]$$

701 and:

$$\mathbb{E} \left[\sum_{a \in [K]} \mathbb{P}[a_t = a] \hat{\ell}_{t,a}^2 \right] = \mathbb{E} \left[\sum_{a \in [K]} \frac{\mathbb{I}[a_t = a] \hat{\ell}_{t,a}^2}{\mathbb{P}[a_t = a]} \right] = \sum_{a \in [K]} \hat{\ell}_{t,a}^2 \leq K$$

702 Substituting these equations into Equation (17) (after taking expectations) gives us:

$$\mathbb{E}[\ln(\psi_{t,0})] \leq -\eta \mathbb{E}[\ell_{t,a_t}] + \eta^2 K/2$$

703 which, upon substitution into Equation (13) gives us:

$$\mathbb{E}[\Delta_t - \Delta_{t+1}] \geq \eta(\mathbb{E}[\ell_{t,a_t}] - \ell_{t,y(x_t)}) - \eta^2 K/2 \quad (18)$$

704 Note that this equation implies that the same equation also holds when the expectation is not implicitly
705 conditional on the state of the algorithm at the start of trial t . Hence, we now drop the assumption that
706 the expectation is conditional on the state of the algorithm at the start of trial t . Summing Equation
707 (18) over all trials $t \in [T]$ and then rearranging gives us:

$$\mathbb{E}[R(y)] \leq \frac{1}{\eta} (\mathbb{E}[\Delta_1] - \mathbb{E}[\Delta_{T+1}]) + \frac{\eta K T}{2} \quad (19)$$

708 Now consider a trial t . For all $v \in \mathcal{B}^\dagger$ let:

$$V_t(v) := \sum_{\mathcal{S} \in 2^{\mathcal{X}}} \mathbb{I}[x_t \in \mathcal{S}] w_{t+1}(\triangleleft(v), \mathcal{S}) + \sum_{\mathcal{S} \in 2^{\mathcal{X}}} \mathbb{I}[x_t \in \mathcal{S}] w_{t+1}(\triangleright(v), \mathcal{S})$$

709 Now take any $j \in [\log(K) - 1] \cup \{0\}$ and let $v := v_{t,j}$. Note that:

$$V_t(v) = \beta_t(\triangleleft(v)) \theta_t(\triangleleft(v)) + \beta_t(\triangleright(v)) \theta_t(\triangleright(v))$$

710 so that by definition of $\pi_t(\triangleleft(v))$ and $\pi_t(\triangleright(v))$ we have:

$$V_t(v) = (\theta_t(\triangleleft(v)) + \theta_t(\triangleright(v))) (\pi_t(\triangleleft(v)) \beta_t(\triangleleft(v)) + \pi_t(\triangleright(v)) \beta_t(\triangleright(v)))$$

711 Without loss of generality assume that $\triangleleft(v) \in \mathcal{P}_t$. Then the above equation implies that:

$$V_t(v) = (\theta_t(\triangleleft(v)) + \theta_t(\triangleright(v))) \frac{\pi_t(\triangleleft(v))\psi_{t,j+1} + \pi_t(\triangleright(v))}{\psi_{t,j}}$$

712 so by definition of $\psi_{t,j}$ we have:

$$V_t(v) = (\theta_t(\triangleleft(v)) + \theta_t(\triangleright(v))) = \sum_{\mathcal{S} \in 2^{\mathcal{X}}} \llbracket x_t \in \mathcal{S} \rrbracket w_t(\triangleleft(v), \mathcal{S}) + \sum_{\mathcal{S} \in 2^{\mathcal{X}}} \llbracket x_t \in \mathcal{S} \rrbracket w_t(\triangleright(v), \mathcal{S})$$

713 Note that this equation trivially holds for all $v \in \mathcal{B}^\dagger \setminus \mathcal{P}_t$ and hence holds for all $v \in \mathcal{B}^\dagger$. Since
 714 for all such v and all \mathcal{S} with $x_t \notin \mathcal{S}$ we have $w_{t+1}(\triangleleft(v), \mathcal{S}) = w_t(\triangleleft(v), \mathcal{S})$ and $w_{t+1}(\triangleright(v), \mathcal{S}) =$
 715 $w_t(\triangleright(v), \mathcal{S})$ we then have:

$$\sum_{\mathcal{S} \in 2^{\mathcal{X}}} w_{t+1}(\triangleleft(v), \mathcal{S}) + \sum_{\mathcal{S} \in 2^{\mathcal{X}}} w_{t+1}(\triangleright(v), \mathcal{S}) = \sum_{\mathcal{S} \in 2^{\mathcal{X}}} w_t(\triangleleft(v), \mathcal{S}) + \sum_{\mathcal{S} \in 2^{\mathcal{X}}} w_t(\triangleright(v), \mathcal{S})$$

716 so, by induction on t we have, for all $t \in [T+1]$, that:

$$\sum_{\mathcal{S} \in 2^{\mathcal{X}}} w_t(\triangleleft(v), \mathcal{S}) + \sum_{\mathcal{S} \in 2^{\mathcal{X}}} w_t(\triangleright(v), \mathcal{S}) = 1$$

717 Hence, for all $v \in \mathcal{B} \setminus r(\mathcal{B})$ and $\mathcal{S} \in 2^{\mathcal{X}}$, we have $w_t(v, \mathcal{S}) \in [0, 1]$. We have now shown that
 718 $\Delta_{T+1} \geq 0$ so that Equation 19 gives us:

$$\mathbb{E}[R(y)] \leq \frac{1}{\eta} \mathbb{E}[\Delta_1] + \frac{\eta KT}{2}$$

719 which, by definition of Δ_1 , gives us the desired result.

720 F.6 Theorem C.2

721 The fact that the weighting w_t is valid is given by the following lemma:

722 **Lemma F.1.** *For all $v \in \mathcal{B}^\dagger$ we have:*

$$\sum_{\mathcal{S} \in 2^{\mathcal{X}}} (w_1(\triangleleft(v), \mathcal{S}) + w_1(\triangleright(v), \mathcal{S})) = 1$$

723 *Proof.* We will show that for all $v \in \mathcal{B}^\dagger$ we have:

$$\sum_{\mathcal{S} \in 2^{\mathcal{X}}} w_1(v, \mathcal{S}) = \frac{1}{2}$$

724 which directly implies the result. So take some arbitrary $v \in \mathcal{B}^\dagger$. Define, for all $t \in [T]$, the sets:

$$\mathcal{X}'_t := \{x_s \mid s \in [t]\} \setminus \{x_1\} \quad \text{and} \quad \mathcal{F}_t := \{0, 1\}^{\mathcal{X}'_t \cup \{x_1\}}$$

725 and for all $x \in \mathcal{X}'_t$, $f \in \mathcal{F}_t$ and $i \in [\log(T)]$, define the quantity:

$$\beta_i(x, f) := \llbracket f(x) \neq f(n(x)) \rrbracket 2^i / T + \llbracket f(x) = f(n(x)) \rrbracket (1 - 2^i / T)$$

726 which is defined since $n(x) \in \mathcal{X}'_t \cup \{x_1\}$. Now fix some $i \in [\log(T)]$. For all $t \in [T-1]$ we have:

$$\sum_{f \in \mathcal{F}_{t+1}} \prod_{x \in \mathcal{X}'_{t+1}} \beta_i(x, f) = \sum_{f \in \mathcal{F}_t} \left(\prod_{x \in \mathcal{X}'_t} \beta_i(x, f) \right) \sum_{f_{(x_{t+1})} \in \{0, 1\}} \beta_i(x_{t+1}, f)$$

727 Given any $f \in \mathcal{F}_t$ we have:

$$\sum_{f_{(x_{t+1})} \in \{0, 1\}} \beta_i(x_{t+1}, f) = \left(1 - \frac{2^i}{T}\right) + \frac{2^i}{T} = 1$$

728 and hence:

$$\sum_{f \in \mathcal{F}_{t+1}} \prod_{x \in \mathcal{X}'_{t+1}} \beta_i(x, f) = \sum_{f \in \mathcal{F}_t} \prod_{x \in \mathcal{X}'_t} \beta_i(x, f)$$

729 Since $\mathcal{X}'_T = \mathcal{X}'$ this implies, by induction, that:

$$\sum_{f \in \mathcal{F}_T} \prod_{x \in \mathcal{X}'} \beta_i(x, f) = \sum_{f \in \mathcal{F}_1} \prod_{x \in \mathcal{X}'_1} \beta_i(x, f) = \sum_{f \in \mathcal{F}_1} \prod_{x \in \emptyset} \beta_i(x, f) = \sum_{f \in \mathcal{F}_1} 1 = |\mathcal{F}_1| = 2 \quad (20)$$

730 Note that we have a bijection $\mathcal{G} : \mathcal{F}_T \rightarrow 2^{\mathcal{X}}$ defined by:

$$\mathcal{G}(f) := \{x \in \mathcal{X} \mid f(x) = 1\} \quad \forall f \in \mathcal{F}_T$$

731 and that for all $(i, f, x) \in [\log(T)] \times \mathcal{F}_T \times \mathcal{X}'$ we have:

$$\beta_i(x, f) = \sigma(x, \mathcal{G}(f))2^i/T + (1 - \sigma(x, \mathcal{G}(f)))(1 - 2^i/T)$$

732 Hence, Equation (20) shows us that for all $i \in [\log(T)]$ we have:

$$\sum_{\mathcal{S} \in 2^{\mathcal{X}}} \prod_{x \in \mathcal{X}'} \left(\sigma(x, \mathcal{S}) \frac{2^i}{T} + (1 - \sigma(x, \mathcal{S})) \left(1 - \frac{2^i}{T}\right) \right) = 2$$

733 This implies that:

$$\sum_{\mathcal{S} \in 2^{\mathcal{X}}} w_1(v, \mathcal{S}) = \frac{1}{2}$$

734 which implies the result. \square

735 Now that we have shown that the weighting w_1 is valid we can utilise Theorem C.1 to prove our
736 regret bound. For any set $\mathcal{S} \in 2^{\mathcal{X}}$ define:

$$\phi(\mathcal{S}) := \sum_{x \in \mathcal{X}'} \sigma(x, \mathcal{S})$$

737 For any $i \in [\log(T)]$ define the function $f_i : [T - 1] \rightarrow \mathbb{R}$ by

$$f_i(c) := \left(1 - \frac{2^i}{T}\right)^{T-1-c} \left(\frac{2^i}{T}\right)^c$$

738 for all $c \in [T - 1]$. Choose any set $\mathcal{S} \in 2^{\mathcal{X}}$ and define:

$$j := \min\{\lceil \log(\phi(\mathcal{S}) + 1) \rceil, \log(T) - 1\}$$

739 If $\phi(\mathcal{S}) \geq T/2$ then $j = \log(T) - 1$ so $2^j/T = 1/2$ and hence:

$$-\ln(f_j(\phi(\mathcal{S}))) = (T - 1) \ln(2) \leq 2\phi(\mathcal{S}) \ln(2) \in \mathcal{O}(\phi(\mathcal{S})) \quad (21)$$

740 Now consider the case in which $\phi(\mathcal{S}) < T/2$. Let $h := 2^j/T$. In this case $2^j/T \leq 1/2$ and hence f_j
741 is monotonic decreasing so since $2^j \geq \phi(\mathcal{S})$ we have:

$$\ln(f_j(\phi(\mathcal{S}))) \geq \ln(f_j(2^j)) = \ln(f_j(Th)) \geq T((1 - h) \ln(1 - h) + h \ln(h)) \geq -Th \ln(e/h)$$

742 so since $\phi(\mathcal{S}) + 1 \geq 2^j/2 = Th/2$ and $h \geq 1/T$ we have:

$$-\ln(f_j(\phi(\mathcal{S}))) \leq 2(\phi(\mathcal{S}) + 1) \ln(eT) \in \mathcal{O}((\phi(\mathcal{S}) + 1) \ln(T)) \quad (22)$$

743 Equations (21) and (22) show us that for all possible values of $\phi(\mathcal{S})$ we have:

$$-\ln(f_j(\phi(\mathcal{S}))) \in \mathcal{O}(\ln(T)(\phi(\mathcal{S}) + 1))$$

744 Noting that for all $v \in \mathcal{B}'$ we have $w_1(v, \mathcal{S}) \geq (1/4 \log(T)) f_j(\phi(\mathcal{S}))$ we have now shown that:

$$-\ln(w_1(v, \mathcal{S})) \in \mathcal{O}(\ln(T)(\phi(\mathcal{S}) + 1)) \quad (23)$$

745 for all $v \in \mathcal{B}'$. As in the statement of Theorem C.1 define, for all $v \in \mathcal{B}$, the set:

$$\mathcal{Q}(v) := \{x \in \mathcal{X} \mid y(x) \in \Downarrow(v)\}$$

746 First note that the graph (with vertex set \mathcal{X}) formed by linking x to $n(x)$ for every $x \in \mathcal{X}'$ is a tree so
747 that $\Phi(y) \geq |\{y(x) \mid x \in \mathcal{X}\}| - 1$. So since for all $v \in \mathcal{B}'$ we have $\mathcal{Q}(v) \neq \emptyset$ if and only if v has a

748 descendent in $\{y(x) \mid x \in \mathcal{X}\}$ and each element of $\{y(x) \mid x \in \mathcal{X}\}$ has $\log(K)$ ancestors in \mathcal{B}' we
 749 have:

$$\sum_{v \in \mathcal{B}'} \llbracket \mathcal{Q}(v) \neq \emptyset \rrbracket \leq \log(K) |\{y(x) \mid x \in \mathcal{X}\}| \leq \log(K) (\Phi(y) + 1) \quad (24)$$

750 Now suppose we have some $x \in \mathcal{X}'$. If $y(x) = y(n(x))$ then for all $v \in \mathcal{B}'$ we have $x, n(x) \in \mathcal{Q}(v)$
 751 or $x, n(x) \notin \mathcal{Q}(v)$ and hence $\sigma(x, \mathcal{Q}(v)) = 0$. On the other hand, if $y(x) \neq y(n(x))$ then for any
 752 $v \in \mathcal{B}'$ with $\sigma(x, \mathcal{Q}(v)) = 1$ we must have that either $x \in \mathcal{Q}(v)$ or $n(x) \in \mathcal{Q}(v)$ so v is an ancestor
 753 of either x or $n(x)$ and hence there can be at most $2 \log(K)$ such v . So in any case we have:

$$\sum_{v \in \mathcal{B}'} \sigma(x, \mathcal{Q}(v)) \leq \llbracket y(x) \neq y(n(x)) \rrbracket 2 \log(K)$$

754 Hence we have:

$$\sum_{v \in \mathcal{B}'} \phi(\mathcal{Q}(v)) = \sum_{x \in \mathcal{X}'} \sum_{v \in \mathcal{B}'} \sigma(x, \mathcal{Q}(v)) \leq 2 \log(K) \Phi(y) \quad (25)$$

755 Equation (23) gives us:

$$- \sum_{v \in \mathcal{B}'} \llbracket \mathcal{Q}(v) \neq \emptyset \rrbracket \ln(w_1(v, \mathcal{Q}(v))) \in \mathcal{O} \left(\ln(T) \sum_{v \in \mathcal{B}'} \phi(\mathcal{Q}(v)) + \ln(T) \sum_{v \in \mathcal{B}'} \llbracket \mathcal{Q}(v) \neq \emptyset \rrbracket \right)$$

756 Substituting in equations (24) and (25) then gives us:

$$- \sum_{v \in \mathcal{B}'} \llbracket \mathcal{Q}(v) \neq \emptyset \rrbracket \ln(w_1(v, \mathcal{Q}(v))) \in \mathcal{O}(\ln(K) \ln(T) \Phi(y))$$

757 so by Theorem C.1 we have:

$$\mathbb{E}[R(y)] \in \mathcal{O} \left(\frac{\eta K T}{2} + \frac{\ln(K) \ln(T) \Phi(y)}{\eta} \right)$$

758 Since $\eta = \rho \sqrt{\ln(K) \ln(T) / K T}$ we obtain the result.

759 F.7 Theorem D.1

760 Define $\lambda' : \mathcal{X} \rightarrow \mathbb{R}_+$ as follows. Given $x \in \mathcal{X}$, if there exists a leaf $u \in \mathcal{J}^*$ with $\gamma(u) = x$ then
 761 $\lambda'(x) = \lambda(u)$. Otherwise $\lambda'(x) = 1$. Given $t \in [T]$ define $\hat{\lambda}_t : \mathcal{Z}_t \rightarrow \mathbb{R}_+$ such that for all $u \in \mathcal{Z}_t$
 762 we have that $\hat{\lambda}_t(u) := \lambda'(\gamma(u))$ if u is a leaf of \mathcal{Z}_t and $\hat{\lambda}_t(u) := 1$ otherwise. For all $t \in [T]$ and
 763 $f : \{x_{t'} \mid t' \in [t]\} \rightarrow \{0, 1\}$ define:

$$\mathcal{N}(f) := \{f' \in \{0, 1\}^{\mathcal{Z}_t} \mid \forall u \in \mathcal{Z}_t^*, f'(u) = f(\gamma(u))\}$$

764 and:

$$\hat{w}(f) := \left(\prod_{t' \in [t]: f(x_{t'})=1} \lambda'(x_{t'}) \right) \prod_{t' \in [t] \setminus \{1\}} (\llbracket f(x_{t'}) \neq f(n(x_{t'})) \rrbracket \epsilon + \llbracket f(x_{t'}) = f(n(x_{t'})) \rrbracket (1 - \epsilon))$$

765 and:

$$\hat{v}(f) := \sum_{f' \in \mathcal{N}(f)} \prod_{u \in \mathcal{Z}_t \setminus r(\mathcal{Z}_t)} \tilde{\tau}_{f'(\uparrow_{\mathcal{Z}_t}(u)), f'(u)}(\mathcal{Z}_t, u, \epsilon) \tilde{\kappa}_{f'(u)}(\hat{\lambda}_t, u)$$

766 We now have the following lemma:

767 **Lemma F.2.** For all $t \in [T]$ and $f : \{x_{t'} \mid t' \in [t]\} \rightarrow \{0, 1\}$ we have:

$$\hat{w}(f) = \hat{v}(f)$$

768 *Proof.* We prove by induction on t . Suppose the result holds for $t = s$ (for some $s \geq 2$). We now
 769 show that it holds for $t = s + 1$ as well. Let f^* be the restriction of f onto the set $\{x_{t'} \mid t' \in [s]\}$.
 770 Let u^* and u' be the unique leaves in \mathcal{Z}_{s+1}^* of which $\gamma(u') = n(x_{s+1})$ and $\gamma(u^*) = x_{s+1}$. By the
 771 construction of \mathcal{Z}_{s+1} these vertices are siblings. Let u'' be the parent (in \mathcal{Z}_{s+1}) of both u^* and u' .
 772 First note that:

$$\llbracket f(x_{s+1}) = 0 \rrbracket + \llbracket f(x_{s+1}) = 1 \rrbracket \lambda'(x_{s+1}) = \tilde{\kappa}_{f(x_{s+1})}(\hat{\lambda}_{s+1}, u^*) \quad (26)$$

773 Since, by the construction of \mathcal{Z}_{s+1} , we have $\gamma(\uparrow_{\mathcal{Z}_{s+1}}(u^*)) = \gamma(u'') = n(x_{s+1})$ we also have that
 774 $d(\uparrow_{\mathcal{Z}_{s+1}}(u^*)) = d(u^*) - 1$ so that, since $\phi_1(\epsilon) = \epsilon$, we have:

$$\llbracket f(x_{s+1}) \neq f(n(x_{s+1})) \rrbracket \epsilon + \llbracket f(x_{s+1}) = f(n(x_{s+1})) \rrbracket (1 - \epsilon) = \tilde{\tau}_{f(n(x_{s+1})), f(x_{s+1})}(\mathcal{Z}_{s+1}, u^*, \epsilon) \quad (27)$$

775 Equations (26) and (27) give us:

$$\hat{w}(f) = \hat{w}(f^*) \tilde{\tau}_{f(n(x_{s+1})), f(x_{s+1})}(\mathcal{Z}_{s+1}, u^*, \epsilon) \tilde{\kappa}_{f(x_{s+1})}(\hat{\lambda}_{s+1}, u^*) \quad (28)$$

776 Now suppose we have some $f' \in \mathcal{N}(f)$. We have $\gamma(u'') = \gamma(u')$ and hence $d(\uparrow_{\mathcal{Z}_{s+1}}(u')) =$
 777 $d(u'') = d(u')$ so since $f'(u') = f(n(x_{s+1}))$ and $\phi_0(\epsilon) = 0$ we have:

$$\tilde{\tau}_{f'(\uparrow_{\mathcal{Z}_{s+1}}(u')), f'(u')}(\mathcal{Z}_{s+1}, u', \epsilon) = \tilde{\tau}_{f'(u''), f'(u')}(\mathcal{Z}_{s+1}, u', \epsilon) = \llbracket f'(u'') = f(n(x_{s+1})) \rrbracket \quad (29)$$

778 Since, by the construction of \mathcal{Z}_{s+1} , we have $\uparrow_{\mathcal{Z}_{s+1}}(u'') = \uparrow_{\mathcal{Z}_s}(u')$ and (as above) we have $d(u'') =$
 779 $d(u')$, we also have:

$$\tilde{\tau}_{f'(\uparrow_{\mathcal{Z}_{s+1}}(u'')), f'(u'')}(\mathcal{Z}_{s+1}, u'', \epsilon) = \tilde{\tau}_{f'(\uparrow_{\mathcal{Z}_s}(u')), f'(u')}(\mathcal{Z}_s, u', \epsilon) \quad (30)$$

780 Since $f'(u^*) = f(x_{s+1})$ and $\uparrow_{\mathcal{Z}_{s+1}}(u^*) = u''$ we have:

$$\tilde{\tau}_{f'(\uparrow_{\mathcal{Z}_{s+1}}(u^*)), f'(u^*)}(\mathcal{Z}_{s+1}, u^*, \epsilon) = \tilde{\tau}_{f'(u''), f(x_{s+1})}(\mathcal{Z}_{s+1}, u^*, \epsilon) \quad (31)$$

781 Now let:

$$\zeta^* := \tilde{\tau}_{f(n(x_{s+1})), f(x_{s+1})}(\mathcal{Z}_{s+1}, u^*, \epsilon) \quad ; \quad \zeta' := \tilde{\tau}_{f'(\uparrow_{\mathcal{Z}_s}(u')), f(n(x_{s+1}))}(\mathcal{Z}_s, u', \epsilon)$$

782 Define:

$$g(f') := \prod_{u \in \mathcal{Z}_s \setminus r(\mathcal{Z}_s)} \tilde{\tau}_{f'(\uparrow_{\mathcal{Z}_s}(u)), f'(u)}(\mathcal{Z}_s, u, \epsilon)$$

783 and:

$$g'(f') := \prod_{u \in \mathcal{Z}_{s+1} \setminus r(\mathcal{Z}_{s+1})} \tilde{\tau}_{f'(\uparrow_{\mathcal{Z}_{s+1}}(u)), f'(u)}(\mathcal{Z}_{s+1}, u, \epsilon)$$

784 Combining equations (29), (30) and (31) gives us:

$$\prod_{u \in \{u^*, u', u''\}} \tilde{\tau}_{f'(\uparrow_{\mathcal{Z}_{s+1}}(u)), f'(u)}(\mathcal{Z}_{s+1}, u, \epsilon) = \llbracket f'(u'') = f(n(x_{s+1})) \rrbracket \zeta^* \zeta' \quad (32)$$

785 For all $u \in \mathcal{Z}_{s+1} \setminus \{u^*, u', u''\}$ we have $\uparrow_{\mathcal{Z}_{s+1}}(u) = \uparrow_{\mathcal{Z}_s}(u)$ so that:

$$\tilde{\tau}_{f'(\uparrow_{\mathcal{Z}_{s+1}}(u)), f'(u)}(\mathcal{Z}_{s+1}, u, \epsilon) = \tilde{\tau}_{f'(\uparrow_{\mathcal{Z}_s}(u)), f'(u)}(\mathcal{Z}_s, u, \epsilon)$$

786 and hence, since $f(n(x_{s+1})) = f'(u')$, we have:

$$g'(f') = \frac{g(f')}{\zeta'} \prod_{u \in \{u^*, u', u''\}} \tilde{\tau}_{f'(\uparrow_{\mathcal{Z}_{s+1}}(u)), f'(u)}(\mathcal{Z}_{s+1}, u, \epsilon)$$

787 Substituting in Equation (32) gives us:

$$g'(f') = g(f') \llbracket f'(u'') = f(n(x_{s+1})) \rrbracket \zeta^* \quad (33)$$

788 We have $\tilde{\kappa}_{f'(u'')}(\hat{\lambda}_{s+1}, u'') = 1$ and for all $u \in \mathcal{Z}_s$ we have $\tilde{\kappa}_{f'(u)}(\hat{\lambda}_{s+1}, u) = \tilde{\kappa}_{f'(u)}(\hat{\lambda}_s, u)$.
 789 Substituting into Equation (33) gives us:

$$g'(f') \prod_{u \in \mathcal{Z}_{s+1}} \tilde{\kappa}_{f'(u)}(\hat{\lambda}_{s+1}, u) = \llbracket f'(u'') = f(n(x_{s+1})) \rrbracket \tilde{\kappa}_{f'(u^*)}(\hat{\lambda}_{s+1}, u^*) \zeta^* g(f') \prod_{u \in \mathcal{Z}_s} \tilde{\kappa}_{f'(u)}(\hat{\lambda}_s, u)$$

790 Summing over all $f' \in \mathcal{N}(f)$ and noting that:

$$\tilde{\kappa}_{f'(r(\mathcal{Z}_{s+1}))}(\hat{\lambda}_{s+1}, r(\mathcal{Z}_{s+1})) = \tilde{\kappa}_{f'(r(\mathcal{Z}_s))}(\hat{\lambda}_s, r(\mathcal{Z}_s)) = 1$$

791 gives us:

$$\hat{v}(f) = \tilde{\kappa}_{f'(u^*)}(\hat{\lambda}_{s+1}, u^*) \zeta^* \hat{v}(f^*)$$

792 By the inductive hypothesis we then have:

$$\hat{v}(f) = \tilde{\kappa}_{f'(u^*)}(\hat{\lambda}_{s+1}, u^*) \zeta^* \hat{w}(f^*)$$

793 which by Equation (28) is equal to $\hat{w}(f)$. We have hence shown that if the inductive hypothesis holds
 794 for $t = s$ then it holds for $t = s + 1$ also. An identical argument shows that the inductive hypothesis
 795 holds for $t = 2$. We have hence shown that the inductive hypothesis holds for all $t \in [T] \setminus \{1\}$. \square

796 We now define a bijection $\mathcal{G} : \{0, 1\}^{\mathcal{X}} \rightarrow 2^{\mathcal{X}}$ by:

$$\mathcal{G}(f) := \{x \in \mathcal{X} \mid f(x) = 1\} \quad \forall f \in \{0, 1\}^{\mathcal{X}}$$

797 Note that for all $f : \mathcal{X} \rightarrow \{0, 1\}$ and all $x \in \mathcal{X}'$ we have:

$$\sigma(x, \mathcal{G}(f))\epsilon + (1 - \sigma(x, \mathcal{G}(f)))(1 - \epsilon) = \llbracket f(x) \neq f(n(x)) \rrbracket \epsilon + \llbracket f(x) = f(n(x)) \rrbracket (1 - \epsilon)$$

798 and:

$$\prod_{x \in \mathcal{G}(f)} \lambda'(x) = \prod_{t' \in [T]: f(x_{t'})=1} \lambda'(x_{t'})$$

799 so that:

$$\tilde{w}(\lambda, \epsilon, \mathcal{G}(f)) = \hat{w}(f)$$

800 and hence, by Lemma F.2, we have:

$$\tilde{w}(\lambda, \epsilon, \mathcal{G}(f)) = \hat{v}(f)$$

801 so that:

$$\sum_{\mathcal{S} \in 2^{\mathcal{X}}} \llbracket \gamma(\hat{u}) \in \mathcal{S} \rrbracket \tilde{w}(\lambda, \epsilon, \mathcal{S}) = \sum_{f \in \{0, 1\}^{\mathcal{X}}} \llbracket f(\gamma(\hat{u})) = 1 \rrbracket \hat{v}(f) \quad (34)$$

802 Since:

$$\bigcup \{ \mathcal{N}(f) \mid f \in \{0, 1\}^{\mathcal{X}}, f(\gamma(\hat{u})) = 1 \} = \{ f' \in \{0, 1\}^{\mathcal{Z}_T} \mid f'(\hat{u}) = 1 \}$$

803 and all sets in this union are disjoint, the right hand side of Equation (34) is equal to:

$$\sum_{f' \in \{0, 1\}^{\mathcal{Z}_T}} \llbracket f'(\hat{u}) = 1 \rrbracket \prod_{u \in \mathcal{Z}_T \setminus r(\mathcal{Z}_T)} \tilde{\tau}_{f'(\uparrow_{\mathcal{Z}_T}(u)), f'(u)}(\mathcal{Z}_T, u, \epsilon) \tilde{\kappa}_{f'(u)}(\hat{\lambda}_T, u) \quad (35)$$

804 Given a vertex $u \in \mathcal{Z}_T \setminus \{r(\mathcal{Z}_T)\}$ define:

$$\mathcal{H}(u) := \downarrow_{\mathcal{Z}_T}(u) \cup \{\uparrow_{\mathcal{Z}_T}(u)\}$$

805 and for all $f : \mathcal{H}(u) \rightarrow \{0, 1\}$ define:

$$\hat{\zeta}(u, f) := \prod_{u' \in \downarrow_{\mathcal{Z}_T}(u)} \tilde{\tau}_{f(\uparrow_{\mathcal{Z}_T}(u')), f(u')}(\mathcal{Z}_T, u', \epsilon)$$

806 **Lemma F.3.** *Given a vertex $u' \in \mathcal{Z}_T \setminus \{r(\mathcal{Z}_T)\}$ and an index $i \in \{0, 1\}$ we have:*

$$\sum_{f \in \{0, 1\}^{\mathcal{H}(u')}} \llbracket f(\uparrow_{\mathcal{Z}_T}(u')) = i \rrbracket \hat{\zeta}(u', f) = 1$$

807 *Proof.* We prove by induction on the height of $\downarrow_{\mathcal{Z}_T}(u')$. If this height is equal to zero then $\mathcal{H}(u') =$
 808 $\{u', \uparrow_{\mathcal{Z}_T}(u')\}$ and for all $f : \mathcal{H}(u') \rightarrow \{0, 1\}$ we have:

$$\hat{\zeta}(u', f) = \tilde{\tau}_{f(\uparrow_{\mathcal{Z}_T}(u')), f(u')}(\mathcal{Z}_T, u', \epsilon)$$

809 Since:

$$\tilde{\tau}_{i,0}(\mathcal{Z}_T, u', \epsilon) + \tilde{\tau}_{i,1}(\mathcal{Z}_T, u', \epsilon) = 1 \quad (36)$$

810 we immediately have the result for the case that the height of $\downarrow_{\mathcal{Z}_T}(u')$ is zero. Now suppose that the
 811 result holds whenever the height of $\downarrow_{\mathcal{Z}_T}(u')$ is equal to j (for some $j \in \mathbb{N}$). We will now show that it
 812 holds whenever the height of $\downarrow_{\mathcal{Z}_T}(u')$ is equal to $j + 1$ which will prove that the result holds always.
 813 By the inductive hypothesis we have, for all $i' \in \{0, 1\}$

$$\sum_{f \in \{0, 1\}^{\mathcal{H}(\triangleleft(u'))}} \llbracket f(u') = i' \rrbracket \hat{\zeta}(\triangleleft(u'), f) = 1$$

814 and

$$\sum_{f \in \{0, 1\}^{\mathcal{H}(\triangleright(u'))}} \llbracket f(u') = i' \rrbracket \hat{\zeta}(\triangleright(u'), f) = 1$$

815 so:

$$\sum_{f \in \{0, 1\}^{\mathcal{H}(u')}} \llbracket f(\uparrow_{\mathcal{Z}_T}(u')) = i \rrbracket \llbracket f(u') = i' \rrbracket \hat{\zeta}(\triangleleft(u'), f) \hat{\zeta}(\triangleright(u'), f) = 1$$

816 and hence:

$$\sum_{f \in \{0, 1\}^{\mathcal{H}(u')}} \llbracket f(\uparrow_{\mathcal{Z}_T}(u')) = i \rrbracket \llbracket f(u') = i' \rrbracket \hat{\zeta}(u', f) = \tilde{\tau}_{i,i'}(\mathcal{Z}_T, u, \epsilon)$$

817 Summing over $i' \in \{0, 1\}$ and noting Equation (36) then shows us the result holds for this case and
 818 hence, by induction, holds always. \square

819 Given $u', u'' \in \mathcal{Z}_T$ with $u'' \in \Downarrow_{\mathcal{Z}_T}(u')$ we define $\hat{\mathcal{H}}(u', u'')$ to be the maximal subtree of \mathcal{Z}_T which
 820 has u' and u'' as leaves. Given, in addition, $f : \hat{\mathcal{H}}(u', u'') \rightarrow \{0, 1\}$ we define:

$$\tilde{\zeta}(u', u'', f) := \prod_{u \in \hat{\mathcal{H}}(u', u'') \setminus \{u'\}} \tilde{\tau}_{f(\uparrow_{\mathcal{Z}_T}(u)), f(u)}(\mathcal{Z}_T, u, \epsilon)$$

821 and:

$$\delta(u', u'') := d(u'') - d(u')$$

822 We now have the following lemma.

823 **Lemma F.4.** Given $u', u'' \in \mathcal{Z}_T$ with $u'' \in \Downarrow_{\mathcal{Z}_T}(u') \setminus \{u'\}$ and indices $i', i'' \in \{0, 1\}$ we have that:

$$\sum_{f \in \{0, 1\}^{\hat{\mathcal{H}}(u', u'')}} \llbracket f(u') = i' \rrbracket \llbracket f(u'') = i'' \rrbracket \tilde{\zeta}(u', u'', f)$$

824 is equal to

$$\llbracket i' \neq i'' \rrbracket \phi_{\delta(u', u'')}(\epsilon) + \llbracket i' = i'' \rrbracket (1 - \phi_{\delta(u', u'')}(\epsilon))$$

825 *Proof.* We prove by induction on the distance from u' to u'' in \mathcal{Z}_T . If this distance is one then we
 826 have $u' = \uparrow_{\mathcal{Z}_T}(u'')$ and $\hat{\mathcal{H}}(u', u'') = \{u', u''\}$ so we have:

$$\sum_{f \in \{0, 1\}^{\hat{\mathcal{H}}(u', u'')}} \llbracket f(u') = i' \rrbracket \llbracket f(u'') = i'' \rrbracket \tilde{\zeta}(u', u'', f) = \tilde{\tau}_{i', i''}(\mathcal{Z}_T, u'', \epsilon)$$

827 which immediately implies that the inductive hypothesis holds in this case. Now suppose that the
 828 inductive hypothesis holds whenever the distance from u' to u'' is j . We now consider the case
 829 that the distance from u' to u'' is $j + 1$. Let u^* be the child of u' that lies in $\hat{\mathcal{H}}(u', u'')$. Without
 830 loss of generality assume that u'' is a descendant of $\triangleleft(u^*)$. Now choose any $i^* \in \{0, 1\}$. Given
 831 $f : \hat{\mathcal{H}}(u', u'') \rightarrow \{0, 1\}$ let:

$$h(i^*, f) = \llbracket f(u') = i' \rrbracket \llbracket f(u'') = i'' \rrbracket \llbracket f(u^*) = i^* \rrbracket$$

832 and let f' and f'' be the restriction of f onto the sets $\hat{\mathcal{H}}(u^*, u'')$ and $\mathcal{H}(\triangleright(u^*))$ respectively. Note that

$$\tilde{\zeta}(u', u'', f) = \tilde{\tau}_{f(u'), f(u^*)}(\mathcal{Z}_T, u^*, \epsilon) \tilde{\zeta}(u^*, u'', f') \hat{\zeta}(\triangleright(u^*), f'')$$

833 By Lemma F.3 and the inductive hypothesis we then have that the quantity:

$$\sum_{f \in \{0, 1\}^{\hat{\mathcal{H}}(u', u'')}} h(i^*, f) \tilde{\zeta}(u', u'', f)$$

834 is equal to the quantity:

$$\tilde{\tau}_{i', i^*}(\mathcal{Z}_T, u^*, \epsilon) (\llbracket i^* \neq i'' \rrbracket \phi_{\delta(u^*, u'')}(\epsilon) + \llbracket i^* = i'' \rrbracket (1 - \phi_{\delta(u^*, u'')}(\epsilon)))$$

835 Summing over $i^* \in \{0, 1\}$ gives us the result. We have hence proved the result in general. \square

836 Suppose we have some $f : \mathcal{J} \rightarrow \{0, 1\}$. Let:

$$\hat{h}(f) = \{f' \in \{0, 1\}^{\mathcal{Z}_T} \mid \forall u \in \mathcal{J}, f'(u) = f(u)\}$$

837 Given $u \in \mathcal{J}$ we have that:

$$\llbracket f(\uparrow_{\mathcal{J}}(u)) \neq f(u) \rrbracket \phi_{\delta(\uparrow_{\mathcal{J}}(u), u)}(\epsilon) + \llbracket f(\uparrow_{\mathcal{J}}(u)) = f(u) \rrbracket (1 - \phi_{\delta(\uparrow_{\mathcal{J}}(u), u)}(\epsilon))$$

838 is equal to $\tilde{\tau}_{f(\uparrow_{\mathcal{J}}(u)), f(u)}(\mathcal{J}, u, \epsilon)$ and hence Lemma F.4 implies that:

$$\sum_{f' \in \hat{\mathcal{H}}(\uparrow_{\mathcal{J}}(u), u)} \llbracket f'(\uparrow_{\mathcal{J}}(u)) = f(\uparrow_{\mathcal{J}}(u)) \rrbracket \llbracket f'(u) = f(u) \rrbracket \tilde{\zeta}(\uparrow_{\mathcal{J}}(u), u, f') = \tilde{\tau}_{f(\uparrow_{\mathcal{J}}(u)), f(u)}(\mathcal{J}, u, \epsilon)$$

839 so since, by the definition of a contraction, the edge sets of the subtrees in $\{\hat{\mathcal{H}}(\uparrow_{\mathcal{J}}(u), u) \mid u \in$
 840 $\mathcal{J} \setminus \{r(\mathcal{J})\}\}$ partition the edge set of \mathcal{Z}_T we have, by definition of $\tilde{\zeta}$, that:

$$\sum_{f' \in \hat{h}(f)} \prod_{u \in \mathcal{Z}_T \setminus \{r(\mathcal{Z}_T)\}} \tilde{\tau}_{f'(\uparrow_{\mathcal{Z}_T}(u)), f'(u)}(\mathcal{Z}_T, u, \epsilon) = \prod_{u \in \mathcal{J} \setminus \{r(\mathcal{J})\}} \tilde{\tau}_{f(\uparrow_{\mathcal{J}}(u)), f(u)}(\mathcal{J}, u, \epsilon)$$

841 Since for all $f' \in \hat{h}(f)$ and for all $u \in \mathcal{Z}_T \setminus \mathcal{J}$ we have $\tilde{\kappa}_{f'(u)}(\hat{\lambda}_T, u) = 1$ we have now shown that
 842 the quantity:

$$\sum_{f' \in \hat{h}(f)} \prod_{u \in \mathcal{Z}_T \setminus \{r(\mathcal{Z}_T)\}} \tilde{\tau}_{f'(\uparrow_{\mathcal{Z}_T}(u)), f'(u)}(\mathcal{Z}_T, u, \epsilon) \tilde{\kappa}_{f'(u)}(\hat{\lambda}_T, u)$$

843 is equal to the quantity:

$$\prod_{u \in \mathcal{J} \setminus \{r(\mathcal{J})\}} \tilde{\tau}_{f(\uparrow_{\mathcal{J}}(u)), f(u)}(\mathcal{J}, u, \epsilon) \tilde{\kappa}_{f(u)}(\hat{\lambda}_T, u)$$

844 Summing over all $f \in \mathcal{F}(\mathcal{J}, \hat{u})$ and noting Equations (34) and (35) gives us the result.

845 F.8 Theorem E.1

846 **Lemma F.5.** *Given $u, u' \in \mathcal{Z}_t$ the algorithm for computing $\nu(u, u')$ is correct.*

847 *Proof.* If $u = u'$ then the proof is trivial. Otherwise we consider the following cases:

- 848 • Consider first the case that $\hat{s} \neq \nabla(s^*)$. Without loss of generality assume $\hat{s} = \triangleleft(s^*)$. Then
 849 we have $u \in \Downarrow(\triangleleft(\xi(s^*)))$ and since $\hat{s}' \neq \triangleleft(s^*)$ we have $u' \notin \Downarrow(\triangleleft(\xi(s^*)))$. Hence $u' \notin \Downarrow(u)$
 850 so $\nu(u, u') = \blacktriangle$ as required.
- 851 • If $u = \xi(s^*)$ then $\hat{s} = \nabla(s^*)$ so either $\hat{s}' = \triangleleft(s^*)$ or $\hat{s}' = \triangleright(s^*)$. In the former case we have
 852 $u' \in \Downarrow(\triangleleft(\xi(s^*))) = \Downarrow(\triangleleft(u))$ so that $\nu(u, u') = \blacktriangleleft$ and similarly in the later case we have
 853 $\nu(u, u') = \blacktriangleright$ as required.
- 854 • If $\hat{s} = \nabla(s^*)$ and $u \neq \xi(s^*)$ we invoke the process. Consider the vertex s at any stage in
 855 the process. By induction we have that if $s \in \mathcal{E}^\bullet$ then $u' \in \Downarrow(\mu'(s))$. This is because if
 856 $s \in \mathcal{E}^\bullet$ then $\mu'(s)$ is an ancestor of $\mu'(\uparrow_{\mathcal{E}}(s))$. This further implies that when $s \neq \hat{s}$ we
 857 have $u' \in \Downarrow(\mu'(\uparrow_{\mathcal{E}}(s)))$. Now suppose that $s \in \mathcal{E}^\circ$ and without loss of generality assume
 858 $s = \triangleleft(\uparrow_{\mathcal{E}}(s))$. Then $u \in \Downarrow(\triangleleft(\xi(\uparrow_{\mathcal{E}}(s))))$ and $\mu'(\uparrow_{\mathcal{E}}(s)) \in \Downarrow(\triangleright(\xi(\uparrow_{\mathcal{E}}(s))))$ so that, since
 859 $u' \in \Downarrow(\mu'(\uparrow_{\mathcal{E}}(s)))$, we have $u' \notin \Downarrow(u)$ and hence $\nu(u, u') = \blacktriangle$ as required. Suppose now
 860 that $s \in \mathcal{E}^\bullet$ and that $u = \xi(s)$. If $\triangleleft(s) \in \mathcal{E}^\bullet$ then we have $\mu'(s) \in \Downarrow(\triangleleft(\xi(s))) = \Downarrow(\triangleleft(u))$
 861 so that, by above, $u' \in \Downarrow(\triangleleft(u))$ and hence $\nu(u, u') = \blacktriangleleft$ as required. Similarly, if $\triangleright(s) \in \mathcal{E}^\bullet$
 862 then $\nu(u, u') = \blacktriangleright$ as required. This completes the proof.

863 □

864 **Lemma F.6.** *The algorithm correctly finds \hat{u} .*

865 *Proof.* By induction on the depth of s we have, for all vertices s in the constructed path, that:

- 866 • If $s \in \mathcal{D}^\circ$ then u_t lies in the maximal spanning tree of \mathcal{Z}_t containing $\mu(s)$ and having
 867 $\uparrow_{\mathcal{J}}(\mu(s))$ as a leaf .
- 868 • If $s \in \mathcal{D}^\bullet$ then u_t lies in the maximal spanning tree of \mathcal{Z}_t with $\uparrow_{\mathcal{J}}(\mu(s))$ and $\mu'(s)$ as
 869 leaves.

870 Let s' be the unique leaf of \mathcal{D} that is on the constructed path. If $s' \in \mathcal{D}^\circ$ then $\mu(s')$ is a leaf of \mathcal{J} and
 871 hence also a leaf of \mathcal{Z}_t . So by above we have that u_t lies in the maximal spanning tree of \mathcal{Z}_t with
 872 $\uparrow_{\mathcal{J}}(\mu(s'))$ and $\mu(s')$ as leaves. If, on the other hand, $s' \in \mathcal{D}^\bullet$ then since s' is a leaf of \mathcal{D} we have that
 873 $\mu(s') = \mu'(s')$ and hence, by above, we have that u_t lies in the maximal spanning tree of \mathcal{Z}_t with
 874 $\uparrow_{\mathcal{J}}(\mu(s'))$ and $\mu(s')$ as leaves. In either case we have $\hat{u} = \mu(s')$ as required. □

875 **Lemma F.7.** *The algorithm correctly finds u^* .*

876 *Proof.* By induction on the depth of s we have, for all vertices s in the constructed path, that:

- 877 • If $s \in \mathcal{E}^\circ$ then $\Gamma_{\mathcal{Z}_t}(u_t, \hat{u})$ lies in $\Downarrow_{\mathcal{Z}_t}(\mu(s))$.

878 • If $s \in \mathcal{E}^\bullet$ then $\Gamma_{\mathcal{Z}_t}(u_t, \hat{u})$ lies in the maximal spanning tree of \mathcal{Z}_t with $\mu(s)$ and $\mu'(s)$ as
879 leaves.

880 Let s' be the unique leaf of \mathcal{E} that is on the constructed path. If $s' \in \mathcal{E}^\circ$ then $\mu(s')$ is a leaf of \mathcal{Z}_t and
881 hence, by above, $\Gamma_{\mathcal{Z}_t}(u_t, \hat{u}) = \mu(s')$ as required. If $s \in \mathcal{E}^\bullet$ then $\mu(s) = \mu'(s)$ and hence, by above,
882 $\Gamma_{\mathcal{Z}_t}(u_t, \hat{u}) = \mu(s')$ as required. \square