# Bayesian Backpropagation Over I-O Functions Rather Than Weights 

David H. Wolpert<br>The Santa Fe Institute<br>1660 Old Pecos Trail<br>Santa Fe, NM 87501


#### Abstract

The conventional Bayesian justification of backprop is that it finds the MAP weight vector. As this paper shows, to find the MAP i-o function instead one must add a correction term to backprop. That term biases one towards i-o functions with small description lengths, and in particular favors (some kinds of) feature-selection, pruning, and weight-sharing.


## 1 INTRODUCTION

In the conventional Bayesian view of backpropagation (BP) (Buntine and Weigend, 1991; Nowlan and Hinton, 1994; MacKay, 1992; Wolpert, 1993), one starts with the "likelihood" conditional distribution $P$ (training set $=\mathbf{t} I$ weight vector $w$ ) and the "prior" distribution $\mathrm{P}(\mathbf{w})$. As an example, in regression one might have a "Gaussian likelihood", $\mathrm{P}(\mathrm{t} \mid \mathbf{w}) \propto$ $\exp \left[-\chi^{2}(w, t)\right] \equiv \Pi_{i} \exp \left[-\left\{\operatorname{net}\left(w, t_{X}(i)\right)-t_{y}(i)\right\}^{2} / 2 \sigma^{2}\right]$ for some constant $\sigma$. ( $t_{X}(i)$ and $t_{Y}(i)$ are the successive input and output values in the training set respectively, and net( $\mathbf{w}$, .) is the function, induced by $w$, taking input neuron values to output neuron values.) As another example, the "weight decay" (Gaussian) prior is $\mathrm{P}(\mathbf{w}) \propto \exp \left(-\alpha\left(\mathbf{w}^{2}\right)\right)$ for some constant $\alpha$.
Bayes' theorem tells us that $\mathrm{P}(\mathbf{w} \mid t) \propto \mathrm{P}(\mathrm{t} \mid \mathrm{w}) \mathrm{P}(\mathbf{w})$. Accordingly, the most probable weight given the data - the "maximum a posteriori" (MAP) $\mathbf{w}$ - is the mode over $\mathbf{w}$ of $\mathrm{P}(\mathrm{t} \mid \mathrm{w}) \mathrm{P}(\mathbf{w})$, which equals the mode over $w$ of the "cost function" $L(w, t) \equiv \ln [P(t \mid w)]+\ln [P(w)]$. So for example with the Gaussian likelihood and weight decay prior, the most probable w given the data is the $w$ minimizing $\chi^{2}(w, t)+\alpha w^{2}$. Accordingly BP with weight decay can be viewed as a scheme for trying to find the function from input neuron values to output neuron values (i-o function) induced by the MAP w.

One peculiar aspect of this justification of weight-decay BP is the fact that rather than the i-o function induced by the most probable weight vector, in practice one would usually prefer to know the most probable i-o function. (In few situations would one care more about a weight vector than about what that weight vector parameterizes.) Unfortunately, the difference between these two i-o functions can be large; in general it is not true that "the most probable output corresponds to the most probable parameter" (Denker and LeCun, 1991).
This paper shows that to find the MAP i-o function rather than the MAP w one adds a "correction term" to conventional BP. That term biases one towards i-o functions with small description lengths, and in particular favors feature-selection, pruning and weight-sharing. In this that term constitutes a theoretical justification for those techniques.

Although cast in terms of neural nets, this paper's analysis applies to any case where convention is to use the MAP value of a parameter encoding Z to estimate the value of Z .

## 2 BACKPROP OVER I-O FUNCTIONS

Assume the net's architecture is fixed, and that weight vectors $\mathbf{w}$ live in a Euclidean vector space $\mathbf{W}$ of dimension IWI. Let $X$ be the set of vectors $x$ which can be loaded on the input neurons, and $O$ the set of vectors 0 which can be read off the output neurons. Assume that the number of elements in $\mathrm{X}(\mid \mathrm{XI})$ is finite. This is always the case in the real world, where measuring devices have finite accuracy, and where the computers used to emulate neural nets are finite state machines. For similar reasons $\mathbf{O}$ is also finite in practice. However for now assume that $O$ is very large and "fine-grained", and approximate it as a Euclidean vector space of dimension IOI. (This assumption usually holds with neural nets, where output values are treated as real-valued vectors.) This assumption will be relaxed later.
Indicate the set of functions taking $X$ to $O$ by $\Phi$. (net(w, .) is an element of $\Phi$.) Any $\phi \in \Phi$ is an ( $|\mathbf{X I} \times| \mathrm{OI}$ )-dimensional Euclidean vector. Accordingly, densities over $\mathbf{W}$ are related to densities over $\Phi$ by the usual rules for transforming densities between |W|-dimensional and ( $\mid \mathrm{XI} \times 1 \mathrm{OI}$ )-dimensional Euclidean vector spaces. There are three cases to consider:

1) $\quad|\mathbf{W}|<|X||O|$. In general, as one varies over all $w$ 's the corresponding i-o functions net( $\mathbf{w},$. ) map out a sub-manifold of $\Phi$ having lower dimension than $\Phi$.
2) $\quad|\mathbf{W}|>|X||O|$. There are an infinite number of $\mathbf{w}$ 's corresponding to each $\phi$.
3) $\quad|\mathbf{W}|=|X||O|$. This is the easiest case to analyze in detail. Accordingly I will deal with it first, deferring discussion of cases one and two until later.
With some abuse of notation, let capital letters indicate random variables and lower case letters indicate values of random variables. So for example w is a value of the weight vector random variable $W$. Use ' $p$ ' to indicate probability densities. So for example $p_{\Phi \mid T}(\phi \mid t)$ is the density of the i-o function random variable $\Phi$, conditioned on the training set random variable T , and evaluated at the values $\Phi=\phi$ and $\mathrm{T}=\mathrm{t}$.
In general, any i-o function not expressible as net( $\mathbf{w}$. .) for some $w$ has zero probability. For the other i-o functions, with $\delta($.$) being the multivariable Dirac delta function,$

$$
\begin{equation*}
p_{\Phi}(\operatorname{net}(w, .))=\int d w^{\prime} p_{w}\left(w^{\prime}\right) \delta\left(\operatorname{net}^{\prime}\left(w^{\prime}, .\right)-\operatorname{net}(w, .)\right) . \tag{1}
\end{equation*}
$$

When the mapping $\Phi=\operatorname{net}(\mathbf{W},$.$) is one-to-one, we can evaluate equation (1) to get$

$$
\begin{equation*}
\mathrm{p}_{\Phi \mid \mathrm{T}}(\operatorname{net}(\mathbf{w}, .) \mid \mathrm{t})=\mathrm{p}_{\mathbf{W} \mid \mathrm{T}}(\mathbf{w} \mid \mathrm{t}) / \mathrm{J}_{\Phi, \mathbf{w}}(\mathbf{w}) \tag{2}
\end{equation*}
$$

where $\mathrm{J}_{\Phi, \mathbf{W}}(\mathbf{w})$ is the Jacobian of the $\mathbf{W} \rightarrow \Phi$ mapping:

$$
\begin{equation*}
\mathrm{J}_{\Phi, \mathbf{w}}(\mathbf{w}) \equiv\left|\operatorname{det}\left[\partial \Phi_{\mathrm{i}} / \partial \mathbf{W}_{\mathbf{j}}\right](\mathbf{w})\right|=\left|\operatorname{det}\left[\partial \operatorname{net}(\mathbf{w},)_{\mathrm{i}} / \partial \mathbf{w}_{\mathrm{j}}\right]\right| . \tag{3}
\end{equation*}
$$

"net( $w,.)_{i}$ " means the $i$ 'th component of the i-o function net $(w,$.$) . "net (w, x)$ " means the vector o mapped by net $(w,$.$) from the input x$, and "net $(w, x)_{k}$ " is the $k$ 'th component of $o$. So the " i " in "net $(w,)_{i}$ " refers to a pair of values $\{x, k\}$. Each matrix value $\partial \phi_{\mathrm{i}} / \partial \mathrm{w}_{\mathrm{j}}$ is the partial derivative of net $(\mathbf{w}, \mathbf{x})_{\mathbf{k}}$ with respect to some weight, for some $x$ and $k . J_{\Phi, w}(\mathbf{w})$ can be rewritten as $\operatorname{det}^{1 / 2}\left[g_{\mathrm{ij}}(w)\right]$, where $\mathrm{g}_{\mathrm{ij}}(\mathrm{w}) \equiv \Sigma_{\mathrm{k}}\left[\left(\partial \phi_{\mathrm{k}} / \partial \mathrm{w}_{\mathrm{i}}\right)\left(\partial \phi_{\mathrm{k}} / \partial \mathrm{w}_{\mathrm{j}}\right)\right]$ is the metric of the $\mathbf{W} \rightarrow \Phi$ mapping. This form of $\mathrm{J}_{\boldsymbol{\Phi}, \mathbf{W}}(\mathbf{w})$ is usually more difficult to evaluate though.

Unfortunately, $\phi=\operatorname{net}(\mathbf{w},$.$) is not one-to-one; where \mathrm{J}_{\Phi, \mathbf{W}}(\mathbf{w}) \neq 0$ the mapping is locally one-to-one, but there are global symmetries which ensure that more than one $w$ corresponds to each $\phi$. (Such symmetries arise from things like permuting the hidden neurons or changing the sign of all weights leading into a hidden neuron - see (Fefferman, 1993) and references therein.) To circumvent this difficulty we must make a pair of assumptions.
To begin, restrict attention to $\mathbf{W}_{\text {inj }}$, those values $\mathbf{w}$ of the variable $\mathbf{W}$ for which the Jacobian is non-zero. This ensures local injectivity of the map between $\mathbf{W}$ and $\Phi$. Given a particular $w \in \mathbf{W}_{\text {inj }}$, let $k$ be the number of $\mathbf{w}^{\prime} \in \mathbf{W}_{\text {inj }}$ such that net $(\mathbf{w},)=.\operatorname{net}\left(\mathbf{w}^{\prime}\right.$, .). (Since $\operatorname{net}(w,)=.\operatorname{net}(w,),. k \geq 1$.) Such a set of $k$ vectors form an equivalence class, $\{w\}$.

The first assumption is that for all $w \in \mathbf{W}_{\text {inj }}$ the size of ( $w$ \} (i.e., $k$ ) is the same. This will be the case if we exclude degenerate $w$ (e.g., w's with all first layer weights set to 0 ). The second assumption is that for all $w^{\prime}$ and $w$ in the same equivalence class, $p_{w \mid D}(w \mid d)=$ $\mathbf{p}_{\mathbf{W} / \mathrm{D}}\left(\mathbf{w}^{\prime} \mid \mathrm{d}\right)$. This is usually the case. (For example, start with $\mathbf{w}^{\prime}$ and relabel hidden neurons to get a new $\mathbf{w} \in\left\{\mathbf{w}^{\prime}\right\}$. If we assume the Gaussian likelihood and prior, then since neither differs for the two w's the weight-posterior is also the same for the two w's.) Given
 usual cost function, $\mathrm{L}(\mathbf{w}, \mathrm{t})$, to find the MAP $\Phi$ BP should minimize $\mathrm{L}^{\prime}(\mathrm{w}, \mathrm{t}) \equiv \mathrm{L}(\mathrm{w}, \mathrm{t})+$ $\ln \left[J_{\Phi, W}(\mathbf{w})\right]$. The $\ln \left[\mathrm{J}_{\Phi, \mathbf{W}}(\mathbf{w})\right]$ term constitutes a correction term to conventional BP.
One should not confuse the correction term with the other quantities in the neural net literature which involve partial derivative matrices. As an example, one way to characterize the "quality" of a local peak $w$ ' of a cost function involves the Hessian of that cost function (Buntine and Weigend, 1991). The correction term doesn't directly concern the validity of such a Hessian-based quality measure. However it does concern the validity of some implementations of such a measure. In particular, the correction term changes the location of the peak $w^{\prime}$. It also suggests that a peak's quality be measured by the Hessian of $L^{\prime}\left(w^{\prime}, t\right)$ with respect to $\phi$, rather than by the Hessian of $L\left(w^{\prime}, t\right)$ with respect to $w$. (As an aside on the subject of Hessians, note that some workers incorrectly use Hessians when they attempt to evaluate quantities like output-variances. See (Wolpert, 1994).)

If we stipulate that the $p_{\Phi \mid T}(\phi \mid t)$ one encounters in the real world is independent of how one chooses to parameterize $\Phi$, then the probability density of our parameter must depend on how it gets mapped to $\Phi$. This is the basis of the correction term. As this suggests, the correction term won't arise if we use non-p $\boldsymbol{p}_{\Phi \mid}(\phi \mid t)$-based estimators, like maximum-likelihood estimators. (This is a basic difference between such estimators and MAP estimators with a uniform prior.) The correction term is also irrelevant if it we use an MAP estimate but $\mathrm{J}_{\Phi, \mathbf{w}}(\mathbf{w})$ is independent of $\mathbf{w}$ (as when net ( $\mathbf{w}$, .) depends linearly on $w$ ). And for nonlinear net( $w,$. ), the correction term has no effect for some non-MAP-based ways to apply Bayesianism to neural nets, like guessing the posterior average $\Phi$ (Neal, 1993):

$$
\begin{equation*}
\mathrm{E}(\Phi \mid \mathrm{t}) \equiv \int \mathrm{d} \phi \mathrm{p}_{\Phi \mid T}(\phi \mid \mathrm{t}) \phi=\int \mathrm{dw} \mathrm{p}_{\mathbf{w} / \mathrm{T}}(\mathrm{w} \mid \mathrm{t}) \operatorname{net}(\mathbf{w}, .) \tag{4}
\end{equation*}
$$

so one can calculate $\mathrm{E}(\Phi \mid \mathrm{t})$ by working in $\mathbf{W}$, without any concern for a correction term. (Loosely speaking, the Jacobian associated with changing integration variables cancels the Jacobian associated with changing the argument of the probability density. A formal derivation - applicable even when $|\mathbf{W}| \neq|\mathbf{X}| \times|O|-$ is in the appendix of (Wolpert, 1994).)
One might think that since it's independent of $t$, the correction term can be absorbed into $\mathbf{p}_{\mathbf{W}}(\mathbf{w})$. Ironically, it is precisely because quantities like $\mathrm{E}(\Phi \mid \mathrm{t})$ aren't affected by the correction term that this is impossible: Absorb the correction term into the prior, giving a new prior $\mathbf{p}^{*} \mathbf{w}(\mathbf{w}) \equiv \mathrm{d} \times \mathrm{p}_{\mathbf{W}}(\mathbf{w}) \times \mathrm{J}_{\Phi, \mathbf{W}}(\mathbf{w})$ (asterisks refers to new densities, and d is a normalization constant). Then $p^{*}{ }_{\Phi \mid T}(\operatorname{net}(w,) \mid t)=.p_{W \mid T}(w \mid t)$. So by redefining what we call the prior we can justify use of conventional uncorrected BP; the (new) MAP $\phi$ corresponds to the $w$ minimizing $L(w, t)$. However such a redefinition changes $E(\Phi \mid t)$ (amongst other
 $\int \mathrm{d} \phi \mathrm{p}_{\Phi \mid \mathrm{T}}(\phi \mid \mathrm{t}) \phi$. So one can either modify BP (by adding in the correction term) and leave $\mathrm{E}(\Phi \mid \mathrm{t})$ alone, or leave BP alone but change $\mathrm{E}(\Phi \mid \mathrm{t})$; one can not leave both unchanged.
Moreover, some procedures involve both prior-based modes and prior-based integrals, and therefore are affected by the correction term no matter how $\mathrm{p}_{\mathbf{W}}(\mathbf{w})$ is redefined. For example, in the evidence procedure (Wolpert, 1993; MacKay, 1992) one fixes the value of a hyperparameter $\Gamma$ (e.g., $\alpha$ from the introduction) to the value $\gamma$ maximizing $\mathrm{p}_{\Gamma \mid \mathrm{T}}\left(\gamma^{\prime} \mid \mathrm{t}\right.$ ). Next one find the value s' maximizing $\mathrm{p}_{\mathrm{SIT}, \Gamma}\left(\mathrm{s}^{\prime} \mid \mathrm{t}, \gamma^{\prime}\right)$ for some variable S . Finally, one guesses the $\phi$ associated with s'. Now it's hard to see why one should use this procedure with $\mathbf{S}=\mathbf{W}$ (as is conventional) rather than with $\mathbf{S}=\boldsymbol{\Phi}$. But with $\mathbf{S}=\Phi$ rather than $\mathbf{W}$, one must factor in the correction term when calculating $\mathrm{p}_{\mathrm{SIT}, \Gamma}\left(\mathrm{s} \mid \mathrm{t}, \gamma^{\prime}\right.$ ), and therefore the guessed $\phi$ is different from when $\mathbf{S}=\mathbf{W}$. If one tries to avoid this change in the guessed $\phi$ by absorbing the correction term into the prior $\mathrm{p}_{\mathbf{W I \Gamma}}(\mathrm{w} \mid \gamma)$, then $\mathrm{p}_{\Gamma \mid \mathrm{T}}(\gamma \mid \mathrm{t})$ - which is given by an integral involving that prior - changes. This in turn changes $\gamma$, and therefore the guessed $\phi$ again is different. So presuming one is more directly interested in $\Phi$ rather than $\mathbf{W}$, one can't avoid having the correction term affect the evidence procedure.
It should be noted that calculating the correction term can be laborious in large nets. One should bear in mind the determinant-evaluation tricks mentioned in (Buntine and Weigend, 1991), as well as others like the identity $\ln \left[J_{\Phi, w}(w)\right]=\operatorname{Tr}\left(\ln \left[\partial \phi_{\mathrm{i}} / \partial \mathbf{w}_{\mathbf{j}}\right]\right) \cong$ $\operatorname{Tr}\left(\ln *\left[\partial \phi_{\mathrm{i}} / \partial \mathbf{w}_{\mathrm{j}}\right]\right)$, where $\ln { }^{*}($.$) is \ln ($.$) evaluated to several orders.$

## 3 EFFECTS OF THE CORRECTION TERM

To illustrate the effects of the correction term, consider a perceptron with a single output neuron, N input neurons and a unary input space: $0=\tanh (\mathbf{w} \cdot \mathbf{x})$, and $\mathbf{x}$ always consist of a single one and N-1 zeroes. For this scenario $\partial \phi_{i} / \partial w_{j}$ is an $N \times N$ diagonal matrix, and $\ln \left[\mathrm{J}_{\Phi, w}(w)\right]=-2 \Sigma_{\mathrm{k}=1}^{\mathrm{N}} \ln \left[\cosh \left(\mathbf{w}_{\mathrm{k}}\right)\right]$. Assume the Gaussian prior and likelihood of the introduction, and for simplicity take $2 \sigma^{2}=1$. Both $L(w, t)$ and $L^{\prime}(w, t)$ are sums of terms each of which only concerns one weight and the corresponding input neuron. Accordingly, it suffices to consider just the $i$ 'th weight and the corresponding input neuron.
Let $\mathbf{x}(i)$ be the input vector which has its 1 in neuron $i$. Let $o_{j}(i)$ be the output of the $j$ 'th of the pairs in the training set with input $\mathbf{x}(\mathrm{i})$, and $\mathrm{m}_{\mathrm{i}}$ the number of such pairs. With $\alpha=0$
(no weight decay), $\mathrm{L}(\mathrm{w}, \mathrm{t})=\chi^{2}(\mathrm{t}, \mathbf{w})$, which is minimized by $\mathbf{w}_{\mathrm{i}}^{\prime}=\tanh ^{-1}\left[\sum_{\mathrm{j}=1}^{\mathrm{m}_{\mathrm{i}}} \mathrm{o}_{\mathrm{j}}(\mathbf{i}) / \mathrm{m}_{\mathrm{j}}\right]$. If we instead try to minimize $\chi^{2}(t, w)+J_{\Phi, w}(w)$ though, then for low enough $m_{i}$ (e.g., $m_{i}$ $=1$ ), we find that there is no minimum. The correction term pushes $w$ away from 0 , and for low enough $m_{i}$ the likelihood isn't strong enough to counteract this push.


Figures 1 through 3: Train using unmodified BP on training set $t$, and feed input $x$ into the resultant net. The horizontal axis gives the output you get. If $t$ and $x$ were still used but training had been with modified BP, the output would have been the value on the vertical axis. In succession, the three figures have $\alpha=.6, .4, .4$, and $m=1,4,1$.


Figure 3.


Figure 4: The horizontal axis is $\left|w_{i}\right|$. The top curve depicts the weight decay regularizer, $\alpha w_{i}{ }^{2}$, and the bottom curve shows that regularizer modified by the correction term. $\alpha=.2$.

When weight-decay is used though, modified BP finds a solution, just like unmodified BP does. Since the correction term "pushes out" $\mathbf{w}$, and since tanh(.) grows with its argument, $\mathrm{a} \phi$ found by modified BP has larger (in magnitude) values of o than does the corresponding $\phi$ found by unmodified BP. In addition, unlike unmodified BP, modified BP has multiple extrema over certain regimes. All of this is illustrated in figures (1) through (3), which graph the value of $o$ resulting from using modified BP with a particular training set t and input value x vs. the value of o resulting from using unmodified BP with t and x . Figure (4) depicts the $w_{i}$-dependences of the weight decay term and of the weight-decay term plus the correction term. (When there's no data, BP searches for minima of those curves.)

Now consider multi-layer nets, possibly with non-unary X. Denote a vector of the compo-
nents of $w$ which lead from the input layer into hidden neuron $K$ by $w_{[K]}$. Let $x^{\prime}$ be the input vector consisting of all 0 's. Then $\partial \tanh \left(\mathbf{w}_{[K]} \cdot \mathbf{x}^{\prime}\right) / \partial \mathbf{w}_{j}=0$ for any $j$, w, and $K$, and for any $w$, there is a row of $\partial \phi_{i} / \partial w_{j}$ which is all zeroes. This in turn means that $J_{\Phi, w}(w)=$ 0 for any $\mathbf{w}$, which means that $\mathbf{W}_{\text {inj }}$ is empty, and $p_{\Phi \mid T}(\phi \mid t)$ is independent of the data $t$. (Intuitively, this problem arises since the o corresponding to $\mathbf{x}^{\prime}$ can't vary with $\mathbf{w}$, and therefore the dimension of $\Phi$ is less than $|\mathbf{W}|$ ) So we must forbid such an all-zeroes $\mathbf{x}^{\prime}$. The easiest way to do this is to require that one input neuron always be on, i.e., introduce a bias unit. An alternative is to redefine $\Phi$ to be the functions from the set $\{\mathrm{X}-(0,0, \ldots, 0)\}$ to O rather than from the set X to O . Another alternative, appropriate when the original X is the set of all input neuron vectors consisting of 0 's and 1 's, is to instead have input neuron values $\in\{z \neq 0,1\}$. (In general $z \neq-1$ though; due to the symmetries of the tanh, for many architectures $z=-1$ means that two rows of $\partial \phi_{i} / \partial w_{j}$ are identical up to an overall sign, which means that $\mathrm{J}_{\Phi, \mathbf{W}}(\mathbf{w})=0$.) This is the solution implicitly assumed from now on.
$\mathrm{J}_{\Phi, \mathbf{W}}(\mathbf{w})$ will be small - and therefore $\mathrm{p}_{\Phi}(\mathrm{net}(\mathbf{w},)$.$) will be large - whenever one can make$ large changes to $w$ without affecting $\phi=\operatorname{net}(w,$.$) much. In other words, p_{\Phi}(\operatorname{net}(w,)$.$) will$ be large whenever we don't need to specify $w$ very accurately. So the correction factor favors those $\mathbf{w}$ which can be expressed with few bits. In other words, the correction factor enforces a sort of automatic MDL (Rissanen, 1986; Nowlan and Hinton, 1994).
More generally, for any multi-layer architecture there are many "singular weights" $\mathbf{w}_{\sin } \notin$ $\mathbf{W}_{\text {inj }}$ such that $\mathrm{J}_{\boldsymbol{\Phi}, \mathbf{W}}\left(\mathbf{w}_{\text {sin }}\right)$ is not just small but equals zero exactly. $\mathrm{p}_{\mathbf{W}}(\mathbf{w})$ must compensate for these singularities, or the peaks of $\mathrm{p}_{\Phi \mid T}(\phi \mid t)$ won't depend on $t$. So we need to have $\mathbf{p}_{\mathbf{w}}(\mathbf{w}) \rightarrow 0$ as $\mathbf{w} \rightarrow \mathbf{w}_{\text {sin }}$. Sometimes this happens automatically. For example often $\mathbf{w}_{\text {sin }}$ includes infinite-valued $\mathbf{w}$ 's, since $\tanh ^{\prime}(\infty)=0$. Because $\mathrm{p}_{\mathbf{w}}(\infty)=0$ for the weightdecay prior, that prior compensates for the infinite-w singularities in the correction term.

For other $\mathbf{w}_{\text {sin }}$ there is no such automatic compensation, and we have to explicitly modify $\mathbf{p}_{\mathbf{W}}(\mathbf{w})$ to avoid singularities. In doing so though it seems reasonable to maintain a "bias" towards the $w_{s i n}$, that $p_{w}(w)$ goes to zero slowly enough so that the values $p_{\Phi}(n e t(w, ~)$. are "enhanced" for $w$ near $w_{s i n}$. Although a full characterization of such enhanced $w$ is not in hand, it's easy to see that they include certain kinds of pruned nets (Hassibi and Stork, 1992), weight-shared nets (Nowlan and Hinton, 1994), and feature-selected nets.

To see that (some kinds of) pruned nets have singular weights, let $\mathrm{w}^{*}$ be a weight vector with a zero-valued weight coming out of hidden neuron K . By (1) $\left.\mathrm{p}_{\Phi}\left(\operatorname{net}^{\left(w^{*},\right.},\right)\right)=$ $\int \mathrm{d}^{\prime} \mathbf{p}_{\mathbf{W}}\left(\mathbf{w}^{\prime}\right) \delta\left(\operatorname{net}\left(\mathbf{w}^{\prime},.\right)\right.$ - net( $\left.\left.\mathbf{w}^{*},.\right)\right)$. Since we can vary the value of each weight $\mathbf{w}^{*}{ }_{i}$ leading into neuron K without affecting net( $\left.\mathbf{w}^{*},.\right)$, the integral diverges. So $\mathbf{w}^{*}$ is singular; removing a hidden neuron results in an enhanced probability. This constitutes an a priori argument in favor of trying to remove hidden neurons during training.

This argument does not apply to weights leading into a hidden neuron; $\mathrm{J}_{\boldsymbol{\Phi}, \mathbf{W}}(\mathbf{w})$ treats weights in different layers differently. This fact suggests that however $\mathbf{p}_{\mathbf{w}}(\mathbf{w})$ compensates for the singularities in $\mathrm{J}_{\Phi, \mathbf{W}}(\mathbf{w})$, weights in different layers should be treated differently by $\mathrm{p}_{\mathbf{W}}(\mathbf{w})$. This is in accord with the advice given in (MacKay, 1992).

To see that some kinds of weight-shared nets have singular weights, let $w$ ' be a weight vector such that for any two hidden neurons K and $\mathrm{K}^{\prime}$ the weight from input neuron i to K equals the weight from $i$ to $\mathrm{K}^{\prime}$, for all input neurons i . In other words, $w$ is such that all hid-
den neurons compute identical functions of $\mathbf{x}$. (For some architectures we'll actually only need a single pair of hidden neurons to be identical.) Usually for such a situation there is a pair of columns of the matrix $\partial \phi_{\mathrm{i}} / \partial \mathbf{w}_{\mathrm{j}}$ which are exactly proportional to one another. (For example, in a 3-2-1 architecture, with $\mathrm{X}=\{\mathrm{z}, 1\}^{3},|\mathbf{W}|=|\mathrm{X}| \times|\mathrm{O}|=8$, and there are four such pairs of columns.) This means that $\mathrm{J}_{\Phi, \mathrm{W}}\left(\mathbf{w}^{\prime}\right)=0 ; \mathbf{w}^{\prime}$ has an enhanced probability, and we have an a priori argument in favor of trying to equate hidden neurons during training.
The argument that feature-selected nets have singular weights is architecture-dependent, and there might be reasonable architectures for which it fails. To illustrate the argument, consider the 3-2-1 architecture. Let $x_{1}(k)$ and $x_{2}(k)$ with $k=\{1,2,3\}$ designate three distinct pairs of input vectors. For each $k$ have $x_{1}(k)$ and $x_{2}(k)$ be identical for all input neurons except neuron $A$, for which they differ. (Note there are four pairs of input vectors with this property, one for each of the four possible patterns over input neurons B and C.) Let $\mathbf{w}^{\prime}$ be a weight vector such that both weights leaving A equal zero. For this situation $\operatorname{net}\left(w^{\prime}, x_{1}(k)\right)=\operatorname{net}\left(w^{\prime}, x_{2}(k)\right)$ for all $k$. In addition $\partial \operatorname{net}\left(w, x_{1}(k)\right) / \partial w_{j}=$ $\partial \operatorname{net}\left(w, x_{2}(k)\right) / \partial w_{j}$ for all weights $w_{j}$ except the two which lead out of $A$. So $k=1$ gives us a pair of rows of the matrix $\partial \phi_{\mathrm{i}} / \partial w_{\mathrm{j}}$ which are identical in all but two entries (one row for $x_{1}(k)$ and one for $\left.x_{2}(k)\right)$. We get another such pair of rows, differing from each other in the exact same two entries, for $k=2$, and yet another pair for $k=3$. So there is a linear combination of these six rows which is all zeroes. This means that $\mathrm{J}_{\Phi, w}\left(\mathbf{w}^{\prime}\right)=0$. This constitutes an a priori argument in favor of trying to remove input neurons during training.

Since it doesn't favor any $\mathrm{p}_{\mathbf{w}}(\mathbf{w})$, the analysis of this paper doesn't favor any $\mathrm{p}_{\Phi}(\phi)$. However when combined with empirical knowledge it suggests certain $\mathrm{p}_{\Phi}(\phi)$. For example, there are functions $\mathrm{g}(\mathrm{w})$ which empirically are known to be good choices for $\mathrm{p}_{\Phi}(\mathrm{net}(\mathrm{w},)$. (e.g., $\left.g(w) \propto \exp \left[\alpha w^{2}\right]\right)$. There are usually problems with such choices of $p_{\Phi}(\phi)$ though. For example, these $g(w)$ usually make more sense as a prior over $W$ than as a prior over $\Phi$, which would imply $\mathbf{p}_{\Phi}($ net $(w,))=.g(w) / J_{\Phi, w}(w)$. Moreover it's empirically true that enhanced $w$ should be favored over other $w$, as advised by the correction term. So it makes sense to choose a compromise between $g(w)$ and $g(w) / J_{\Phi, w}(w)$. An example is $p_{\Phi}(\phi) \propto$ $g(w) /\left[\lambda_{1}+\tanh \left(\lambda_{2} \times \mathrm{J}_{\Phi, w}(\mathbf{w})\right)\right]$ for two hyperparameters $\lambda_{1}>0$ and $\lambda_{2}>0$.

## 4 BEYOND THE CASE OF BACKPROP WITH |W| = |X||O|

When O does not approximate a Euclidean vector space, elements of $\Phi$ have probabilities rather than probability densities, and $\mathrm{P}(\phi \mid t)=\int \mathrm{d}_{\mathrm{w}} \mathrm{p}_{\mathbf{W I T}}(\mathrm{w} \mid \mathrm{t}) \delta($ net $(\mathrm{w},),. \phi),(\delta(.,$.$) being$ a Kronecker delta function). Moreover, if O is a Euclidean vector space but $\mathbf{W}|>|X|| O \mid$, then again one must evaluate a difficult integral; $\Phi=n e t(W,$.$) is not one-to-one so one$ must use equation (1) rather than (2). Fortunately these two situations are relatively rare.

The final case to consider is $|\mathbf{W}|<|X||O|$ (see section two). Let $\mathbf{S}(\mathbf{W})$ be the surface in $\Phi$ which is the image (under net( $\mathbf{W},$.$) ) of \mathbf{W}$. For all $\phi \mathrm{p}_{\Phi}(\phi)$ is either zero (when $\phi \notin \mathbf{S}(\mathbf{W})$ ) or infinite (when $\phi \in S(\mathbf{W})$ ). So as conventionally defined, "MAP $\phi$ " is not meaningful.

One way to deal with this case is to embed the net in a larger net, where that larger net's output is relatively insensitive to the values of the newly added weights. An alternative that is applicable when $|W|$ / $|O|$ is an integer is to reduce X by removing "uninteresting" $x$ 's. A third alternative is to consider surface densities over $S(W), p_{S(W)}(\phi)$, instead of vol-
ume densities over $\Phi, \mathrm{p}_{\Phi}(\phi)$. Such surface densities are given by equation (2), if one uses the metric form of $\mathrm{J}_{\Phi, \mathbf{W}}(\mathbf{w})$. (Buntine has emphasized that the Jacobian form is not even defined for $|\mathbf{W}|<|X||O|$, since $\partial \phi_{\mathrm{i}} / \partial \mathbf{w}_{\mathrm{j}}$ is not square then (personal communication).)
As an aside, note that restricting $\mathrm{p}_{\Phi}(\phi)$ to $\mathrm{S}(\mathbf{W})$ is an example of the common theoretical assumption that "target functions" come from a pre-chosen "concept class". In practice such an assumption is usually ludicrous - whenever it is made there is an implicit hope that it constitutes a valid approximation to a more reasonable $\mathrm{p}_{\Phi}(\phi)$.

When decision theory is incorporated into Bayesian analysis, only rarely does it advise us to evaluate an MAP quantity (i.e., use BP). Instead Bayesian decision theory usually advises us to evaluate quantities like $\mathrm{E}(\Phi \mid \mathrm{t})$ (Wolpert, 1994). Just as it does for the use of MAP estimators, the analysis of this paper has implications for the use of such $\mathrm{E}(\Phi \mid \mathrm{t})$ estimators. In particular, one way to evaluate $E(\Phi \mid t)=\int d w p_{w / T}(w \mid t) \operatorname{net}(w,$.$) is to$ expand net( $\mathbf{w},$. ) to low order and then approximate $p_{W_{T}}(w \mid t)$ as a sum of Gaussians (Buntine and Weigend, 1991). Equation (4) suggests that instead we write $\mathrm{E}(\Phi \mid \mathrm{t}$ ) as $\int \mathrm{d} \phi \mathrm{p}_{\Phi \mid \mathrm{T}}(\phi \mid \mathrm{t}) \phi$ and approximate $\mathrm{p}_{\Phi \mid \mathrm{T}}(\phi \mid \mathrm{t})$ as a sum of Gaussians. Since fewer approximations are used (no low order expansion of net( $w,$.$) ), this might be more accurate.$

## Acknowledgements

Thanks to David Rosen and Wray Buntine for stimulating discussion, and to TXN and the SFI for funding. This paper is a condensed version of (Wolpert 1994).

## References

Buntine, W., Weigend, A. (1991). Bayesian back-propagation. Complex Systems, 5, p. 603.
Denker, J., LeCun, Y. (1991). Transforming neural-net output levels to probability distributions. In Neural Information Processing Systems 3, R. Lippman et al. (Eds).
Fefferman, C. (1993). Reconstructing a neural net from its output. Sarnoff Research Center TR 93-01.

Hassibi, B., and Stork, D. (1992). Second order derivatives for network pruning: optimal brain surgeon. Ricoh Tech Report CRC-TR-9214.
MacKay, D. (1992). Bayesian Interpolation, and A Practical Framework for Backpropagation Networks. Neural Computation, 4, pp. 415 and 448.
Neal, R. (1993). Bayesian learning via stochastic dynamics. In Neural Information Processing Systems 5, S. Hanson et al. (Eds). Morgan Kaufmann.

Nowlan, S., and Hinton, G. (1994). Simplifying Neural Networks by Soft Weight-Sharing. In Theories of Induction: Proceedings of the SFI/CNLS Workshop on Formal Approaches to Supervised Learning, D. Wolpert (Ed.). Addison-Wesley, to appear.
Rissanen, J. (1986). Stochastic complexity and modeling. Ann. Stat., 14, p. 1080.
Wolpert, D. (1993). On the use of evidence in neural networks. In Neural Information Processing Systems 5, S. Hanson et al. (Eds). Morgan-Kauffman.

Wolpert, D. (1994). Bayesian back-propagation over i-o functions rather than weights. SFI tech. report. ftp'able from archive.cis.ohio-state.edu, as pub/neuroprose/wolpert.nips.93.Z.

