Integrating First-Order Logic Programs And Connectionist Systems

A Constructive Approach

Project Thesis by Andreas Witzel

*1979-12-09 in Freiburg, Germany
Matriculation Number: 3065439

Supervisors:

Prof. Dr. rer. nat. habil. Steffen Hölldobler
Dipl.-Inf. Sebastian Bader
Dr. Pascal Hitzler

April 2005
Department of Computer Science
Technische Universität Dresden, Germany

Task

Student:

Name:	Andreas Witzel
Date and Place of Birth:	1979-12-09, Freiburg, Germany
Matriculation Number:	3065439

Topic:

The integration of the paradigms of logic programs and connectionist systems is desirable because of their contrasting advantages and disadvantages. So far, algorithms for transforming logic programs into standard architecture connectionist systems exist only for the case of propositional logic. For first-order logic programs, there exist only non-constructive existence proofs.

The aim of this thesis is to find concrete methods for transforming covered logic programs into standard architecture connectionist systems.

Goals:

- Reviewing the embedding of the T_P operator into the real numbers
- Approximating the embedding with an easy-to-handle function
- Constructing standard architecture connectionist systems computing or approximating that function
- Discussing consequences, problems, alternatives, and possible extensions

Abstract

Significant advances have recently been made concerning the integration of symbolic knowledge representation with connectionist systems (also called artificial neural networks). However, while the integration with propositional paradigms has resulted in applicable systems, the case of first-order knowledge representation has so far hardly proceeded beyond theoretical studies which prove the existence of connectionist systems for approximating first-order logic programs up to any chosen precision. Advances were hindered severely by the lack of concrete algorithms for obtaining the approximating networks which were known to exist: the corresponding proofs are non-constructive in that they do not yield concrete methods for building the systems. In this paper, we will make the required advance and show how to obtain the structure and the parameters for different kinds of connectionist systems approximating covered logic programs.

Contents

1 Introduction

Logic programs have been studied thoroughly in computer science and artificial intelligence and are well understood. They are human-readable, they basically consist of logic formulae, and there are well-founded mathematical theories defining exactly the meaning of a logic program. Logic programs thus constitute one of the most prominent paradigms for knowledge representation and reasoning. But there is also a major drawback: Logic programming is unsuitable for certain learning tasks, in particular in the full first-order case.

On the other hand, for connectionist systems — also called artificial neural networks — there are established and rather simple training or learning algorithms. But it is hard to manually construct a connectionist system with a desired behaviour, and even harder to find a declarative interpretation of what a given connectionist system does. Connectionist systems perform very well in certain settings, but in general we do not understand why or how.

Thus, logic programs and connectionist systems have contrasting advantages and disadvantages. It would be desirable to integrate both approaches in order to combine their respective advantages while avoiding the disadvantages. We could then train a connectionist system to fulfil a certain task, and afterwards translate it into a logic program in order to understand it or to prove that it meets a given specification. Or we might write a logic program and turn it into a connectionist system which could then be optimised using a training algorithm.

Main challenges for the integration of symbolic and connectionist knowledge thus center around the questions (1) how to extract logical knowledge from trained connectionist systems, and (2) how to encode symbolic knowledge within such systems. We find it natural to start with (2), as extraction methods should easily follow from successful methods for encoding.

For propositional logic programs, encodings into connectionist systems like [11] led immediately to applicable algorithms. Corresponding learning paradigms have been developed [7, 6] and applied to real settings.

For the first-order logic case, however, the situation is much more difficult, as laid out in [4]. Concrete translations, as in [3, 2], yield nonstandard network architectures. For standard architectures, previous work has only established non-constructive proofs showing the existence of connectionist systems which approximate given logic program with arbitrary precision [12, 9]. Thus the implementation of first-order integrated systems was impossible up to this point.

In this paper, we will give concrete methods to compute the structure and the parameters of connectionist systems approximating certain logic programs using established standard architectures.

First, in Section 2, we will give a short introduction to logic programs and connectionist systems. We also review the standard technique for bridging the symbolic world of logic programs with the real-numbers-based world of connectionist systems, namely the embedding of the single-step operator, which carries the meaning of a logic program, into the real numbers as established for this purpose in [12]. In Section 3, we will then approximate the resulting real function by a piecewise constant function in a controlled manner, which is an important simplifying step for establishing our re-

sults. We will then construct connectionist systems for computing or approximating this function, using sigmoidal activation functions in Section 4 and radial basis function (RBF) architecture in Section 5. Section 6 will conclude the paper with a short discussion of some open problems and possibilities for future work.

2 Preliminaries

In this section, we shortly review the basic notions needed from logic pogramming and connectionist systems. Main references for background reading are [13] and [14], respectively. We also review the embedding of T_P into the real numbers as used in [12, 9], and on which our approach is based.

2.1 Logic Programs

A *logic program* over some first-order language $\mathcal{L}$ is a set of (implicitly universally quantified) *clauses* of the form $A \leftarrow L_1 \wedge \cdots \wedge L_n$, where $n \in \mathbb{N}$ may differ for each clause, A is an *atom* in $\mathcal{L}$ with variables from a set $\mathcal{V}$, and the L_i are *literals* in $\mathcal{L}$, that is, atoms or negated atoms. A is called the *head* of the clause, the L_i are called *body literals*, and their conjunction $L_1 \wedge \cdots \wedge L_n$ is called the *body* of the clause. As an abbreviation, we will sometimes replace $L_1 \wedge \cdots \wedge L_n$ by body and write $A \leftarrow \text{body}$. If $n = 0$, A is called a *fact*. A clause is *ground* if it does not contain any variables. *Local variables* are those variables occurring in some body but not in the corresponding head. A logic program is *covered* if none of the clauses contain local variables.

Example 2.1. *The following is a covered logic program which will serve as our running example. The intended meaning of the clauses is given to the right.*

$e(0).$	*% 0 is even*
$e(s(X)) \leftarrow \neg e(X)$	*% the successor $s(X)$*
	% of a non-even X is even

The *Herbrand universe* $\mathcal{U}_P$ is the set of all ground terms of $\mathcal{L}$, the *Herbrand base* $\mathcal{B}_P$ is the set of all ground atoms. A *ground instance* of a literal or a clause is obtained by replacing all variables by terms from $\mathcal{U}_P$. For a logic program P, $\mathcal{G}(P)$ is the set of all ground instances of clauses from P.

A *level mapping* is a function $\|\cdot\| : \mathcal{B}_P \to \mathbb{N} \setminus \{0\}$. In this paper, we require level mappings to be injective, in which case they can be thought of as enumerations of $\mathcal{B}_P$. The *level* of an atom A is denoted by $\|A\|$. The level of a literal is that of the corresponding atom.

A logic program P is *acyclic with respect to a level mapping* $\|\cdot\|$ if for all clauses $A \leftarrow L_1 \wedge \cdots \wedge L_n \in \mathcal{G}(P)$ we have that $\|A\| > \|L_i\|$ for $1 \leq i \leq n$. A logic program is called *acyclic* if there exists such a level mapping. All acyclic programs are also covered under our standing condition that level mappings are injective, and provided that function symbols are present, i.e. $\mathcal{B}_P$ is infinite. Indeed the case when $\mathcal{B}_P$ is finite is of limited interest to us as it reduces to a propositional setting as studied in [11, 7].

Example 2.2. *For the program from Example 2.1, we have:*

$$\mathcal{U}_P = \{0, s(0), s^2(0), \dots\}$$
$$\mathcal{B}_P = \{e(0), e(s(0)), e(s^2(0)), \dots\}$$
$$\mathcal{G}(P) = \left\{ \begin{array}{c} e(0). \\ e(s(0)) \leftarrow \neg e(0). \\ e(s(s(0))) \leftarrow \neg e(s(0)). \\ \vdots \end{array} \right\}$$

We will use the common abbreviation $s^n(0)$ for the n-fold application of s to 0. With $\|e(s^n(0))\| := n+1$, we find that P is acyclic. We will always use this level mapping in our running example.

A *(Herbrand) interpretation* is a subset I of $\mathcal{B}_P$. Those atoms A with $A \in I$ are said to be *true*, or to *hold*, under I (in symbols: $I \models A$), those with $A \notin I$ are said to be *false*, or to *not hold*, under I (in symbols: $I \not\models A$). $\mathcal{I}_P = 2^{\mathcal{B}_P}$ is the set of all interpretations.

An interpretation I is a *(Herbrand) model* of a logic program P (in symbols: $I \models P$) if I is a model for each clause $A \leftarrow \text{body} \in \mathcal{G}(P)$ in the usual sense. That is, if of all body literals I contains exactly those which are not negated (i.e. $I \models \text{body}$), then I must also contain the head.

Example 2.3. *Consider these three Herbrand interpretations for P from Example 2.1:*

$$I_1 = \{e(0), e(s(0))\}$$
$$I_2 = \{e(0), e(s^3(0)), e(s^4(0)), e(s^5(0)), \dots\}$$
$$I_3 = \mathcal{B}_P$$

$I_1 \not\models P$ since $e(s^3(0)) \leftarrow \neg e(s^2(0)) \in \mathcal{G}(P)$ and $e(s^2(0)) \notin I_1$, but $e(s^3(0)) \notin I_1$. Both I_2 and I_3 are models for P.

The *single-step operator* $T_P : \mathcal{I}_P \to \mathcal{I}_P$ maps an interpretation I to the set of exactly those atoms A for which there is a clause $A \leftarrow \text{body} \in \mathcal{G}(P)$ with $I \models \text{body}$. The operator T_P captures the semantics of P as the Herbrand models of the latter are exactly the pre-fixed points of the former, i.e. those interpretations I with $T_P(I) \subseteq I$. For logic programming purposes it is usually preferrable to consider fixed points of T_P, instead of pre-fixed points, as the intended meaning of programs. These fixed points are called *supported models* of the program [1]. The well-known stable models [8], for example, are always supported. In Example 2.1, $\{e(0), e(s^2(0)), e(s^4(0)), \dots\}$ is supported (and stable), while B_P is a model but not supported.

Example 2.4. *For P from Example 2.1 and I_1, I_2 from Example 2.3, we get the following by successive application (i.e. iteration) of T_P:*

$$I_1 \overset{T_P}{\mapsto} I_2 \overset{T_P}{\mapsto} \{e(0), e(s^2(0)), e(s^3(0))\} \overset{T_P}{\mapsto} \dots$$
$$\overset{T_P}{\mapsto} \{e(0), e(s^2(0)), \dots, e(s^{2n}(0)), e(s^{2n+1}(0))\} \overset{T_P}{\mapsto} \dots$$

For a certain class of programs, the process of iterating T_P can be shown to converge[1] to the unique supported Herbrand model of the program, which in this case is the model describing the semantics of the program [10]. This class is described by the fact that T_P is a contraction with respect to a certain metric. A more intuitive description remains to be found, but at least all acyclic programs[2] are contained in this class. That is, given some acyclic program P, we can find its unique supported Herbrand model by iterating T_P and computing a limit. In Example 2.4 for instance, the iteration converges in this sense to $\{e(0), e(s^2(0)), e(s^4(0)), \dots\}$, which is the unique supported model of the program.

2.2 Connectionist Systems

A *connectionist system* — or *artificial neural network* — is a complex network of simple computational units, also called *nodes* or *neurons*, which accumulate real numbers from their inputs and send a real number to their output. Each unit's output is *connected* to other units' inputs with a certain real-numbered *weight*. Those units without incoming connections are called *input units* or *input neurons*, those without outgoing ones are called *output units* or *output neurons*.

We will deal with layered feed-forward networks, i.e. networks without cycles where the outputs of units in one layer are only connected to the inputs of units in the next layer, as shown in Figure 1. The first and last layers contain the input and output units respectively, the remaining layers are called *hidden layers*.

Each unit has an *input function* which merges its inputs into one single input using the connections' weights, and an *activation function* which then computes the output. If a unit has inputs $x_1, \dots, x_n$ with weights $w_1, \dots, w_n$, then the *weighted sum* input function is $\sum_{i=1}^{n} x_i w_i$. A locally receptive *distance* input function is given by $\sqrt{\sum_{i=1}^{n} (x_i - w_i)^2}$. In the case of one single input, this is equivalent to $|x_1 - w_1|$.

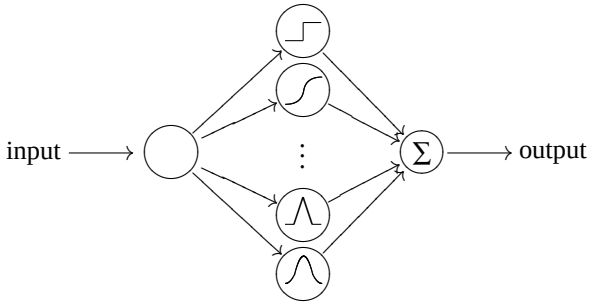

Figure 1: A simple 3-layered feed-forward connectionist system, with different activation functions depicted in the hidden layer.

[1] Convergence in this case is convergence with respect to the Cantor topology on $\mathcal{I}_P$, or equivalently, with respect to a natural underlying metric. For further details, see [10], where also a general class of programs, called Φ-*accessible programs*, is described, for which iterating T_P always converges in this sense.

[2] In this case the level mapping does not need to be injective.

Example 2.5. *The connectionist system shown in Figure 2 classifies its input by outputting* 0 *for inputs* ≤ 5, 1 *for inputs between* 5 *and* 7, *and* 2 *for inputs* > 7.

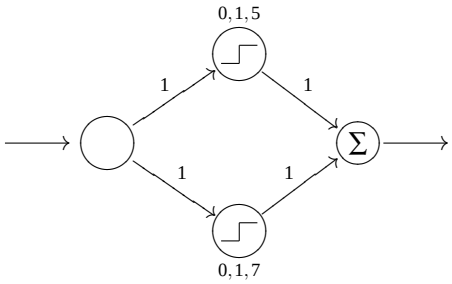

Figure 2: The classifying connectionist system from Example 2.5. The *weights* of the connections are denoted at the arrows representing them. In the hidden layer, we use *weighted sum input functions* and *step activation functions* whose parameters l, h, m are denoted above or below the units. They output l if the merged input is $\leq m$, and $l+h$ if it is $> m$. The output unit simply sums up its weighted inputs, which is indicated by the $\sum$ sign. You could view it as a unit with weighted sum input function and identity output function. Note that, in a real application, the units shown in the hidden layer might be built from several simpler units.

Given a connectionist system, it is hard to read off any meaning beyond the obvious purely mathematical meaning, for example some kind of symbolic or logical interpretation. Vice versa, given some description of desired behaviour, it is not straightforward to design a corresponding connectionist system. However, one of the main advantages of connectionist systems is that there exist good learning algorithms which can be used to train or adapt existing systems. A prominent method is the *backpropagation algorithm*. It changes the network's parameters and performs a gradient descent in order to minimize the deviation from a given desired output. Thus, it requires the network to be differentiable, and that is why our goal will always be to eventually obtain networks with differentiable activation functions. Since we will only use the differentiable input functions described above, and since function composition preserves differentiability, the whole network will then be differentiable, as desired. It is also clear that a derivative of 0 does not help much in a gradient descent, so we would like to use functions whose derivative is never 0.

2.3 Embedding T_P in $\mathbb{R}$

Connectionist systems propagate real numbers, while single-step operators operate on interpretations, i.e. subsets of $\mathcal{B}_P$. Thus, we need to bridge the gap between the real-valued and the symbolic setting. We follow the idea laid out first in [12], and further developed in [9], for embedding $\mathcal{I}_P$ into $\mathbb{R}$. For this purpose, we define $R : \mathcal{I}_P \to \mathbb{R}$ as $R(I) := \sum_{A \in I} b^{-\|A\|}$ for some base $b \geq 3$. Note that R is injective. We will abbreviate $R(\{A\})$ by $R(A)$ for singleton interpretations. Furthermore, we will append a subscript b to floating point numbers with base b.

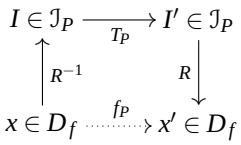

Figure 3: Relations between T_P and f_P

Example 2.6. *With the level mapping defined in Example 2.2, we obtain the following values for the embedding of the interpretations from Example 2.3:*

$$R(I_1) = 0.1100000\ldots_b$$
$$R(I_2) = 0.1001111\ldots_b$$
$$R(I_3) = 0.1111111\ldots_b$$

We define the *embedding of T_P in $\mathbb{R}$* , $f_P : D_f \to D_f$ with $D_f := \{R(I) | I \in \mathcal{I}_P\}$, as $f_P(x) := R(T_P(R^{-1}(x)))$. Figure 3 illustrates the relations between T_P, R, and f_P, while Figure 4 shows the graph of f_P for the program introduced in Example 2.1.

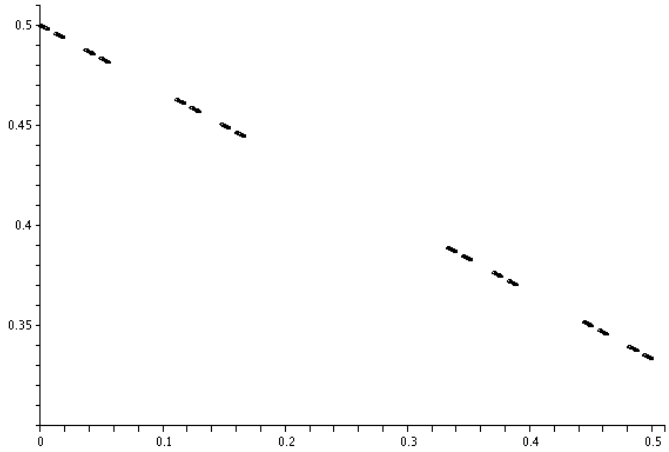

Figure 4: The graph of the embedded T_P-operator from Example 2.1, using base 3 for the embedding. Each point is shown as a small circle. In general, the graph will be more complex and the points will not be on a straight line, so in the following we will apply our general methods and not try to exploit the simplicity of our running example.

3 Constructing Piecewise Constant Functions

In the following, we assume P to be a covered program with bijective level mapping $\|\cdot\|$ which is, along with its inverse $\|\cdot\|^{-1}$, effectively computable. As already mentioned, we also assume that $\mathcal{B}_P$ is infinite. However, our approach will also work for the finite case with minor modifications. Furthermore, R and f_P denote the embeddings with base b as defined above.

3.1 Approximating one Application of T_P

In this section we will show how to construct a ground subprogram approximating a given program. That is, we will construct a subset P_l of the ground program $\mathcal{G}(P)$ such that the associated consequence operator T_{P_l} approximates T_P up to a given accuracy.

Definition 3.1. For all $l \in \mathbb{N}$, the set of *atoms of level less than or equal to l* is defined as

$$\mathcal{A}_l := \{A \in \mathcal{B}_P \mid \|A\| \leq l\}.$$

Furthermore, we define the *instance of P up to level l* as

$$P_l := \left\{A \leftarrow \text{body} \in \mathcal{G}(P) \middle| A \in \mathcal{A}_l\right\}.$$

Since the level mappings are required to be enumerations, we know that $\mathcal{A}_l$ is finite. Furthermore, it is also effectively computable, due to the required computability of $\|\cdot\|^{-1}$. It is clear from the definition that P_l is ground and finite, and again, can be computed effectively.

Example 3.2. *In our running example, we have:*

$$\begin{aligned} \mathcal{A}_1 &= \{e(0)\} \\ \mathcal{A}_2 &= \{e(0), e(s(0))\} \\ \mathcal{A}_3 &= \{e(0), e(s(0)), e(s^2(0))\} \end{aligned} \qquad P_3 = \left\{ \begin{array}{c} e(0). \\ e(s(0)) \leftarrow \neg e(0). \\ e(s^2(0)) \leftarrow \neg e(s(0)). \end{array} \right\}$$

Definition 3.3. For all $l \in \mathbb{N}$, the *greatest relevant input level* with respect to l is

$$\hat{l} := \max\left\{\|L\| \middle| L \text{ is a body literal of some clause in } P_l\right\}.$$

Obviously, we can compute $\hat{l}$ easily, since P_l is ground and finite.

Example 3.4. *In our running example, we have* $\hat{3} = 2$*, since the maximum level of all body literals occurring in* P_3 *is* $\|\neg e(s(0))\| = 2$.

The following lemma establishes a connection between the consequence operators of some ground subprogram P_k and the original program P.

Lemma 3.5. For all $l, k \in \mathbb{N}$, $k \geq l$, and $I, J \in \mathcal{I}_P$, we have that $T_{P_k}(I)$ and $T_P(J)$ agree on $\mathcal{A}_l$ if I and J agree on $\mathcal{A}_{\hat{l}}$, i.e.

$$I \cap \mathcal{A}_{\hat{l}} = J \cap \mathcal{A}_{\hat{l}} \quad \text{implies} \quad T_{P_k}(I) \cap \mathcal{A}_l = T_P(J) \cap \mathcal{A}_l.$$

Proof. This follows simply from the fact that I and J agree on $\mathcal{A}_{\hat{l}}$, and that T_{P_k} contains all those clauses relating atoms from $\mathcal{A}_{\hat{l}}$ and $\mathcal{A}_l$. Taking this into account we find that T_P and T_{P_k} agree on $\mathcal{A}_l$. □

Definition 3.6. The *greatest relevant output level* with respect to some arbitrary $\varepsilon > 0$ is

$$\begin{aligned} o_\varepsilon &:= \min\left\{ n \in \mathbb{N} \,\middle|\, \sum_{\|A\|>n} R(A) < \varepsilon \right\} \\ &= \min\left\{ n \in \mathbb{N} \,\middle|\, n > -\frac{\ln(b-1)\varepsilon}{\ln b} \right\} \end{aligned}$$

The following theorem connects the embedded consequence operator of the subprogram P_{o_ε} with the desired error bound ε. It constitutes the basis for later approximations using connectionist systems.

Theorem 3.7. For all $\varepsilon > 0$, we have that

$$\left| f_P(x) - f_{P_{o_\varepsilon}}(x) \right| < \varepsilon \quad \text{for all } x \in D_f.$$

Proof. Let $x \in D_f$ be given. From Lemma 3.5, we know that $T_{P_{o_\varepsilon}}(R^{-1}(x)) = R^{-1}(f_{P_{o_\varepsilon}}(x))$ agrees with $T_P(R^{-1}(x)) = R^{-1}(f_P(x))$ on all atoms of level $\leq o_\varepsilon$. Thus, $f_{P_{o_\varepsilon}}(x)$ and $f_P(x)$ agree on the first o_ε digits. So the maximum deviation occurs if all later digits are 0 in one case and 1 in the other. In that case, the difference is $\sum_{\|A\|>n} R(A)$, which is $< \varepsilon$ by definition of o_ε. □

Example 3.8. *In our running example, for* $\varepsilon = 0.02$ *and* $b = 3$ *we have:*

$$o_{0.02} = \min\left\{ n \in \mathbb{N} \,\middle|\, n > -\frac{\ln(2 \cdot 0.02)}{\ln 3} \right\} = \min\left\{ n \in \mathbb{N} \,\middle|\, n > 2.93 \right\} = 3$$

Thus, f_{P_3} *approximates* f_P *up to a maximum error of* 0.02.

3.2 Iterating the Approximation

Now we know that one application of $f_{P_{o_\varepsilon}}$ approximates f_P up to ε. But what will happen if we try to approximate several iterations of f_P? In general, $\hat{o}_\varepsilon$ might be greater than o_ε, that is, the required input precision might be greater than the resulting output precision. In that case, we lose precision with each iteration. So in order to achieve a given output precision after a certain number of steps, we increase our overall precision such that we can afford losing some of it. Since the precision might decrease with each step, we can only guarantee a certain precision for a given maximum number of iterations.

Theorem 3.9. For all $l, n \in \mathbb{N}$, we can effectively compute $l^{(n)}$ such that for all $I \in \mathfrak{I}_P$, $m \leq n$, and $k \geq l^{(n)}$:

$$T^m_{P_k}(I) \text{ agrees with } T^m_P(I) \text{ on } \mathcal{A}_l.$$

Proof. By induction on n. Let $l \in \mathbb{N}$ be given.

base $n = 0$**:** Obviously, $T^0_{P_k}(I) = I = T^0_P(I)$. We set $l^{(0)} := l$.

step $n \rightsquigarrow n+1$**:** By induction hypothesis, we can find $l^{(n)}$ such that for all $I \in \mathfrak{I}_P$, $m \leq n$, and $k \geq l^{(n)}$, $T^m_{P_k}(I)$ agrees with $T^m_P(I)$ on $A_{\hat{l}}$. With $l^{(n+1)} := \max\{l, l^{(n)}\}$, we then have for all $I \in \mathfrak{I}_P$, $m \leq n$, and $k \geq l^{(n+1)}$:

$$\begin{aligned} & T^m_{P_k}(I) \text{ agrees with } T^m_P(I) \text{ on } A_{\hat{l}} && (k \geq l^{(n)}) \\ \Rightarrow \quad & T^{m+1}_{P_k}(I) \text{ agrees with } T^{m+1}_P(I) \text{ on } \mathcal{A}_l && (3.5) \end{aligned}$$

$T^0_{P_k}(I) = I = T^0_P(I)$ completes the Induction Step.

□

It follows that for all $\varepsilon > 0$, we can effectively compute $o^{(n)}_\varepsilon$ such that $\left| f^n_P(x) - f^n_{P_{o^{(n)}_\varepsilon}}(x) \right| < \varepsilon$ for all $x \in D_f$.

This result may not seem completely satisfying. If we want to iterate our approximation, we have to know in advance how many steps we will need at most. Of course, we could choose a very large maximum number of iterations, but then the instance of P up to the corresponding level might become very large. But in the general case, we might not be interested in that many iterations anyway, since T_P does not necessarily converge.

For acyclic programs, however, T_P is guaranteed to converge, and additionally we can prove that we do not lose precision in the application of T_{P_l}. Due to the acyclicity of P we have $\hat{l} < l$, and hence, with respect to $\mathcal{A}_l$, we obtain the same result after n iterations of T_{P_l} as we would obtain after n iterations of T_P. Thus we can approximate the fixed point of T_P by iterating T_{P_l}. To put it formally, we have that $T^n_{P_l}(I)$ agrees with $T^n_P(I)$ on $\mathcal{A}_l$ for acylic P and all $n \in \mathbb{N}$. Thus, in this case we find that $|f^n_P(x) - f^n_{P_{o_\varepsilon}}(x)| < \varepsilon$ for all $x \in D_f$ and all $n \in \mathbb{N}$.

3.3 Simplifying the Domain

Now we have gathered all information and methods necessary to approximate f_P and iterations of f_P. It remains to simplify the domain of the approximation so that we can regard the approximation as a piecewise constant function. We do this by extending D_f to some larger set D_l.

The idea is as follows. Since only input atoms of level $\leq \hat{l}$ play a role in P_l, we have that all $x \in D_f$ which differ only after the $\hat{l}$-th digit are mapped to the same value by f_{P_l}. So we have ranges $[x, x'] \subseteq \mathbb{R}$ of fixed length such that all elements of $[x, x'] \cap D_f$ are mapped to the same value.

Example 3.10. *In our running example, one of the ranges $[x, x']$ is given by the following values for the endpoints, since we have $\hat{l} = 2$:*

$$\begin{aligned} x &= 0.\underbrace{01}_{=}000000\ldots_b \\ x' &= 0.\overbrace{01}111111\ldots_b \end{aligned}$$

Obviously, there are $2^{\hat{l}}$ such ranges, each of length $\sum_{\|A\|>\hat{l}} R(A)$. So we can extend f_{P_l} to a function $\hat{f}_{P_l}$ which has a domain consisting of $2^{\hat{l}}$ disjoint and connected ranges and is constant on each of these ranges. Additionally, the minimum distance between two ranges is greater than or equal to the length of the ranges. We formalize these results in the following.

Definition 3.11. An ordered enumeration of all left borders $d_{l,i}$ $(0 \leq i < 2^{\hat{l}})$ of the intervals constituting D_l can be computed as

$$d_{l,i} := \sum_{j=1}^{\hat{l}} \left(\begin{cases} b^{-j} & \text{if } \left\lfloor \frac{i}{\hat{l}-j+1} \right\rfloor \mod 2 = 1 \\ 0 & \text{otherwise} \end{cases} \right).$$

Each of the intervals has length

$$\lambda_l := \sum_{\|A\|>\hat{l}} R(A) = \frac{1}{(b-1)\cdot b^{\hat{l}}}.$$

We now define

$$D_l := \bigcup_{i=0}^{2^{\hat{l}}-1} D_{l,i} \qquad \text{with } D_{l,i} := [d_{l,i}; d_{l,i} + \lambda_l].$$

Thus, D_l consists of $2^{\hat{l}}$ pieces of equal length.

Example 3.12. *In our running example, we have:*

$$d_{3,0} = 0.00_3 \qquad d_{3,1} = 0.01_3 \qquad d_{3,2} = 0.10_3 \qquad d_{3,3} = 0.11_3$$

$$\lambda_3 = 0.001111\ldots_3 = \frac{1}{2\cdot 3^2} = \frac{1}{18} \qquad D_3 = \bigcup_{i=0}^{2^2-1} D_{3,i}$$

$$D_{3,0} = [0.000000\ldots_3; 0.001111\ldots_3] \qquad D_{3,2} = [0.100000\ldots_3; 0.101111\ldots_3]$$
$$D_{3,1} = [0.010000\ldots_3; 0.011111\ldots_3] \qquad D_{3,3} = [0.110000\ldots_3; 0.111111\ldots_3]$$

The corresponding decimal values are given in Example 4.1.

Lemma 3.13. For all $l \in \mathbb{N}$, D_l is an extension of D_f, i.e.

$$D_l \supseteq D_f$$

Proof. Let $l \in \mathbb{N}$ and $x \in D_f$. Then there is a $d_{l,i}$ which agrees with x on its $\hat{l}$ digits. But $D_{l,i}$ contains all numbers which agree with $d_{l,i}$ on its $\hat{l}$ digits, thus $x \in D_{l,i} \subseteq D_l$. □

Lemma 3.14. For all $l \in \mathbb{N}$, the connected parts of D_l do not overlap and the space between one part and the next is at least as wide as the parts themselves.

Proof. The minimum distance between two parts occurs when the left endpoints differ only in the last, i.e. $\hat{l}$-th, digit. In that case, the distance between these endpoints is $b^{-\hat{l}}$, which is $\geq 2\cdot\lambda_l$ since $b \geq 3$. □

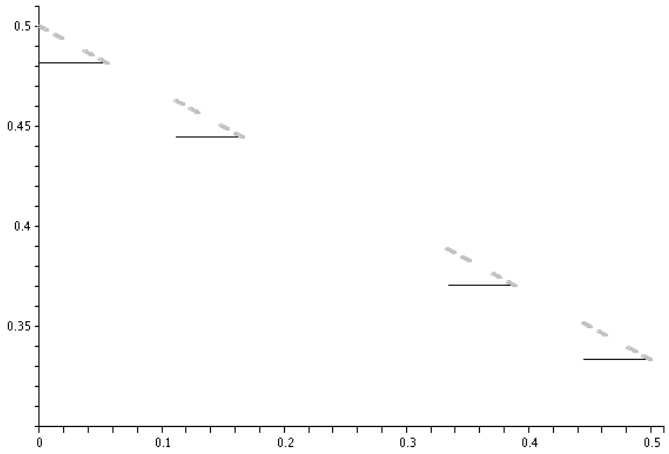

Figure 5: The graph of $\hat{f}_{P_3}$ for our running example; f_P is shown as grey circles.

Lemma 3.15. For all $l \in \mathbb{N}$ and $0 \leq i < 2^{\hat{l}}$, f_{P_l} is constant on $D_{l,i} \cap D_f$.

Proof. All atoms in bodies of clauses of P_l are of level $\leq \hat{l}$. Thus, T_{P_l} regards only those atoms of level $\leq \hat{l}$, i.e. T_{P_l} is constant for all interpretations which agree on these atoms. This means that f_{P_l} is constant for all x that agree on the first $\hat{l}$ digits, which holds for all $x \in D_{l,i} \cap D_f$. □

Definition 3.16. The *extension of* f_{P_l} *to* D_l , $\hat{f}_{P_l} : D_l \to D_f$, is defined as

$$\hat{f}_{P_l}(x) := f_{P_l}(d_{l,i}) \quad \text{for } x \in D_{l,i}.$$

From the results above, it follows that $\hat{f}_{P_l}$ is well-defined.

Example 3.17. *In our running example, we have:*

$$\hat{f}_{P_3}(x) = \begin{cases} 0.111_3 = R\left(\left\{e(0), e(s(0)), e(s^2(0))\right\}\right) & \text{for } x \in D_{3,0} \\ 0.110_3 = R\left(\{e(0), e(s(0))\}\right) & \text{for } x \in D_{3,1} \\ 0.101_3 = R\left(\left\{e(0), e(s^2(0))\right\}\right) & \text{for } x \in D_{3,2} \\ 0.100_3 = R\left(\{e(0)\}\right) & \text{for } x \in D_{3,3} \end{cases}$$

The decimal values are given in Example 4.1, and the corresponding graph is shown in Figure 5. Note also that, no matter where we start, iterated application of $\hat{f}_{P_3}$ *will always end up yielding* $0.101_3 = R\left(\left\{e(0), e(s^2(0))\right\}\right)$*, which is exactly the embedding of the fixed point of* T_P *restricted to* $\mathcal{A}_3$.

Now we have simplified the domain of the approximated embedded single-step operator such that we can regard it as a function consisting of a finite number of equally long constant pieces with gaps at least as wide as their length.

In the following, we will construct connectionist systems which either compute this function exactly or approximate it up to a given, arbitrarily small error. In the latter case we are facing the problem that the two errors might add up to an error which is greater than the desired maximum error. But this is easily taken care of by dividing the desired maximum overall error into one error ε' for $f_{P_{o_{\varepsilon'}}}$ and another error ε'' for the constructed connectionist system.

4 Constructing Sigmoidal Feed-Forward Networks

We will continue our exhibition by considering some arbitrary piecewise constant function g with n equally long pieces, which we want to approximate using connectionist systems. Our $\hat{f}_{P_l}$ is a special case of such a function, so our results will be applicable to it. The simple intention with using g is to save indices and get less complicated expressions. So in the following, let $g : D \to \mathbb{R}$ be given by

$$D := \bigcup_{i=0}^{n-1} [a_i, c_i], \quad c_i = a_i + b, \quad c_i < a_{i+1},$$
$$g(x) := y_i \text{ for } x \in [a_i, c_i].$$

Example 4.1. *In our running example, we have $n = 4$ and the obvious values from Examples 3.12 and 3.17, which translate to the following decimal values:*

$$b = \lambda_3 \approx 0.05556$$

$a_0 = d_{3,0} = 0$	$c_0 = d_{3,0} + \lambda_3 \approx 0.05556$	$y_0 = \hat{f}_{P_3}(d_{3,0}) \approx 0.48148$
$a_1 = d_{3,1} \approx 0.11111$	$c_1 = d_{3,1} + \lambda_3 \approx 0.16667$	$y_1 = \hat{f}_{P_3}(d_{3,1}) \approx 0.44444$
$a_2 = d_{3,2} \approx 0.33333$	$c_2 = d_{3,2} + \lambda_3 \approx 0.38889$	$y_2 = \hat{f}_{P_3}(d_{3,2}) \approx 0.37037$
$a_3 = d_{3,3} \approx 0.44444$	$c_3 = d_{3,3} + \lambda_3 = 0.5$	$y_3 = \hat{f}_{P_3}(d_{3,3}) \approx 0.33333$

When we construct our connectionist systems, we are only interested in the values they yield for inputs in D. We do not care about the values for inputs outside of D since such inputs are guaranteed not to be possible embeddings of interpretations, i.e. in our setting they do not carry any symbolic meaning which could be translated back to $\mathfrak{I}_P$.

We will proceed in two steps. First, we will compute g exactly using a connectionist system with step activation functions. Then, we will replace each step function by a corresponding sigmoidal function similar enough to the former so as to guarantee that a given maximum error for the whole system is not exceeded.

4.1 Step Activation Functions

We will now construct a 3-layered feed-forward network with weighted sum input functions, where each of the units in the hidden layer computes the following step function:

$$s_{l,h,m}(x) := \begin{cases} l & \text{if } x \leq m \\ l+h & \text{otherwise.} \end{cases}$$

As an abbreviation, we will use $s_i(x) := s_{l_i,h_i,m_i}(x)$ for $0 \leq i < n-1$. We want the output to agree with g on its domain, that is, we want $\sum_{i=0}^{n-2} s_i(x) = g(x)$ for all $x \in D$.

An intuitive construction is depicted in Figure 6. For n pieces, we use $n-1$ steps. We put one step in the middle between each two neighbouring pieces, then obviously the height of that step must be the height difference between these two pieces.

It remains to specify values for the left arms of the step functions. All left arms should add up to the height of the first piece. So we can choose that height divided by $n-1$ for each left arm. Now we have specified all s_i completely:

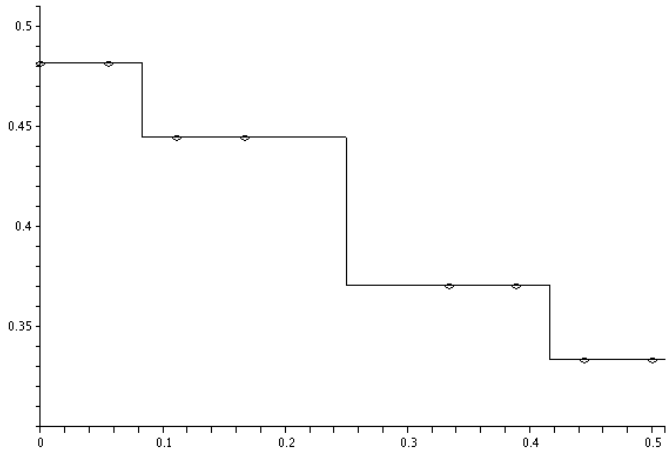

Figure 6: The graph of the function we want to compute using a connectionist system with step activation functions. The circles denote the endpoints of the constant pieces of $\hat{f}_{P_3}$.

Definition 4.2. For $0 \leq i < n-1$,

$$l_i := \frac{y_0}{n-1} \qquad h_i := -y_i + y_{i+1} \qquad m_i := \frac{1}{2}(c_i + a_{i+1})$$

Example 4.3. *For our running example, we get a connectionist system with* 3 *units in the hidden layer. The system is depicted in Figure 7. It computes exactly the function whose graph is shown in Figure 6.*

Theorem 4.4.

$$\sum_{i=0}^{n-2} s_i(x) = g(x) \quad \text{for all } x \in D.$$

Proof. Let $x \in [a_j, c_j]$. Then

$$\begin{aligned}\sum_{i=0}^{n-2} s_i(x) &= \sum_{i=0}^{j-1}(l_i + h_i) + \sum_{i=j}^{n-2} l_i = \sum_{i=0}^{n-2} l_i + \sum_{i=0}^{j-1} h_i \\ &= y_0 + \sum_{i=0}^{j-1}(-y_i + y_{i+1}) = y_j = g(x).\end{aligned}$$

□

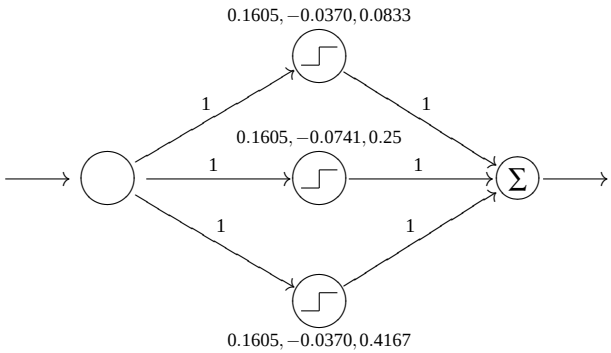

Figure 7: The step function connectionist system for our running example with the approximate values for the parameters l, h, m denoted above or below each respective unit.

4.2 Sigmoidal Activation Functions

Instead of step activation functions, standard network architectures use sigmoidal activation functions, which can be thought of as approximations of step functions. The reason for this is that standard training algorithms like backpropagation require differentiable activation functions.

In order to accomodate this, we will now approximate each step function s_i by a sigmoidal function σ_i:

$$\sigma_i(x) := \sigma_{l_i,h_i,m_i,z_i}(x) := l_i + \frac{h_i}{1+e^{-z_i(x-m_i)}}.$$

Note that l_i, h_i, m_i are the same as for the step functions. The error of the i-th sigmoidal is

$$\delta_i(x) := |\sigma_i(x) - s_i(x)|.$$

An analysis of this function leads to the following results (illustrated in Figure 8): For all $x \neq m_i$ we have

$$\lim_{z_i \to \infty} \sigma_i(x) = s_i(x);$$

since both functions are symmetric, we find for all $z_i, \Delta x$,

$$\delta_i(m_i - \Delta x) = \delta_i(m_i + \Delta x);$$

and furthermore, for all z_i, x, x' with $|x' - m_i| > |x - m_i|$,

$$\delta_i(x') < \delta_i(x).$$

Theorem 4.5. For all $\varepsilon > 0$ we can find z_i ($0 \leq i < n-1$) such that

$$\left| \sum_{i=0}^{n-2} \sigma_i(x) - g(x) \right| < \varepsilon.$$

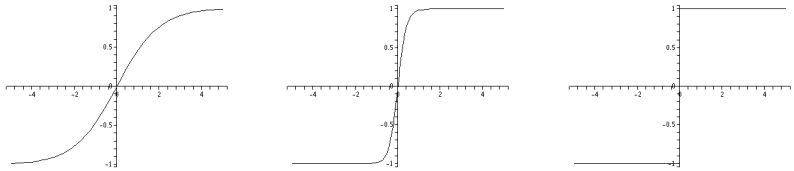

Figure 8: With increasing z, $\sigma_{l,h,m,z}$ gets arbitrarily close to $s_{l,h,m}$ everywhere but at m. The difference between $\sigma_{l,h,m,z}$ and $s_{l,h,m}$ is symmetric to m and decreases with increasing distance from m. Shown here are $\sigma_{-1,2,0,1}, \sigma_{-1,2,0,5}, s_{-1,2,0}$.

Proof. In the worst case, the respective errors of the σ_i add up in the sum. Thus we allow a maximum error of $\varepsilon' := \frac{\varepsilon}{n-1}$ for each σ_i. With all previous results, it only remains to choose the z_i big enough to guarantee that at those $x \in D$ which are closest to m_i (i.e. c_i and a_{i+1}, which are equally close), σ_i approximates s_i up to ε', that is

$$\left[\delta_i(c_i) = \right] \delta_i(a_{i+1}) \overset{!}{<} \varepsilon'.$$

Resolving this we get the following condition for the z_i:

$$z_i > \begin{cases} -\infty & \text{if } |h_i| \leq \varepsilon' \\ -\frac{\ln \varepsilon' - \ln(|h_i| - \varepsilon')}{a_{i+1} - m_i} & \text{otherwise} \end{cases}$$

for $0 \leq i < n-1$. This completes the proof. □

Example 4.6. *We will approximate the network shown in Figure 7 by one with sigmoidal activation functions in the hidden layer. Choosing an approximation accuracy of* $\varepsilon = 0.01$*, we obtain a maximum error of* $\varepsilon' = \frac{0.01}{3}$ *for each sigmoidal unit. With the condition formulated in Theorem 4.5, we can find corresponding values for the* z_i*. The resulting network is shown in Figure 9.*

Note that, since our previous network computes f_{P_3} *which is an approximation of* f_P *up to a maximum error of* 0.02*, we now have a network approximating* f_P *up to a maximum error of* $0.02 + 0.01 = 0.03$*.*

Figure 10 shows the graph of the function computed by this network.

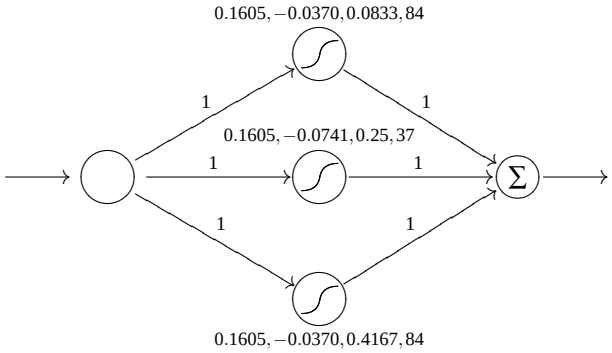

Figure 9: The sigmoidal connectionist system for our running example with the approximate values for the parameters l, h, m, z denoted above or below each respective unit. For each z, the smallest acceptable natural number was chosen.

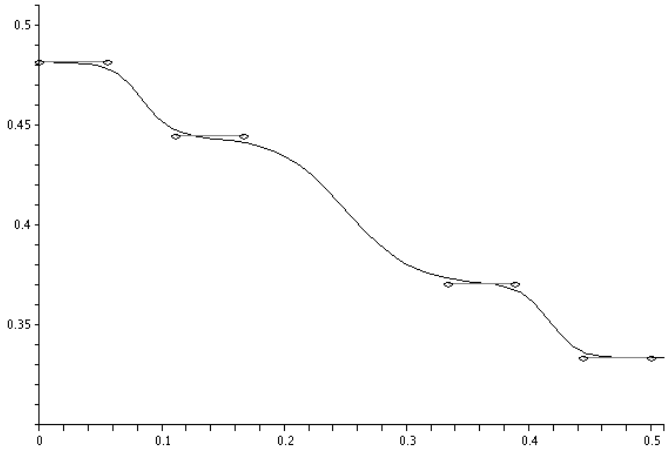

Figure 10: The graph of the function computed by the sigmoidal network of our running example. The constant pieces of $\hat{f}_{P_3}$ are shown in grey with circles at the endpoints.

5 Constructing RBF Networks

In the following section, we will show how to construct *Radial Basis Function Networks* (RBF Networks). For a more detailed introduction to this type of network we refer to [14]. As in the previous section, we will take a two-step approach and first discuss triangular activation functions. We will then extend the results to so-called raised-cosine radial basis functions. We will also briefly discuss how an existing network can be refined incrementally to lower the error bound. The notation is the same as in the previous section. We will again assume that g is a piecewise constant function, this time with the additional requirement that the gaps between the pieces are $\geq$ the length of the pieces (which we proved to hold for $\hat{f}_{P_l}$), i.e. $c_i + b \leq a_{i+1}$ for $0 \leq i < n$.

5.1 Triangular Activation Functions

We will construct an RBF network with distance input function, where each of the units in the hidden layer computes a triangular function $t_{w,h,m}$:

$$t_{w,h,m}(x) := \begin{cases} h \cdot \left(1 - \frac{|x-m|}{w}\right) & \text{if } |x-m| < w \\ 0 & \text{otherwise} \end{cases}$$

Since the triangular functions are locally receptive, that is, they are $\neq 0$ only on the open range $(m-w, m+w)$, we can handle each constant piece separately and represent it as a sum of two triangles, as illustrated in Figure 11.

For each interval $[a_i, c_i]$ (with $c_i = a_i + b$), we define

$$t_i(x) := t_{b,y_i,a_i}(x), \qquad t'_i(x) := t_{b,y_i,c_i}(x).$$

Thus, for each constant piece we get two triangles summing up to that constant piece, i.e. for $0 \leq i < n$ and $x \in [a_i, c_i]$ we have $t_i(x) + t'_i(x) = y_i$, as illustrated in Figure 11.

The requirement we made for the gap between two constant pieces guarantees that the triangles do not interfere with other pieces.

Example 5.1. *For our running example, we get the network depicted in Figure 12.*

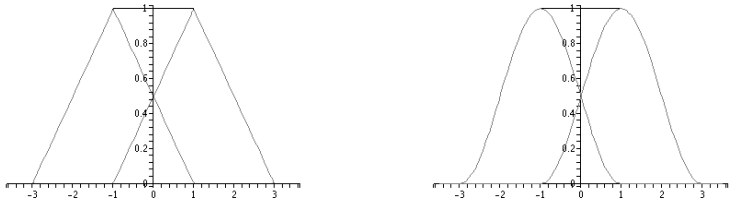

Figure 11: A constant piece can be obtained as the sum of two triangles or two raised-cosine functions.

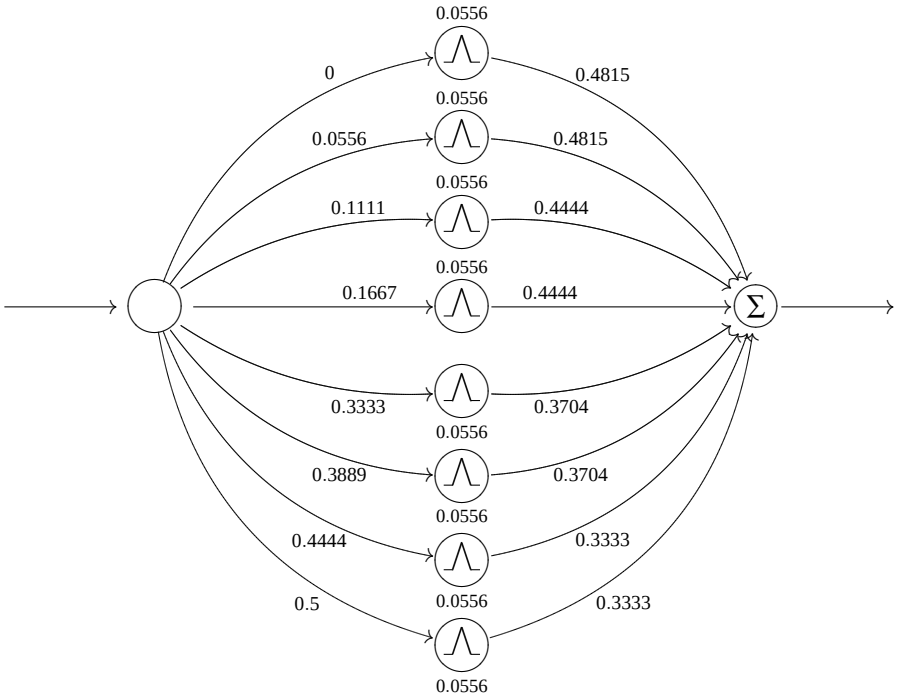

Figure 12: The RBF network for our running example. Since we have distance input functions, the approximate values for m are denoted as weights at the connections to the hidden layer units. The hidden layer units then compute triangles of height 1 with the value for w denoted above or below each respective unit. The results are then scaled in height using the values for h denoted as weights at the connections to the output sum unit. The graph of the function computed by the raised-cosine version of this network is shown in Figure 13.

Theorem 5.2.

$$\sum_{i=0}^{n-1} \big(t_i(x) + t_i'(x)\big) = g(x) \quad \text{for all } x \in D.$$

Proof. This equality follows directly from the fact that the two triangles add up to a constant piece of the required height, and furthermore, that they do not interfere with other constant pieces as mentioned above. □

5.2 Raised-Cosine Activation Functions

As in the previous section, standard radial basis function network architectures use differentiable activation functions. For our purposes, we will replace the triangular functions t_i and t_i' by raised-cosine functions τ_i and τ_i', respectively, of the following

form:

$$\tau_{w,h,m}(x) := \begin{cases} \frac{h}{2} \cdot \left(1 + \cos\left(\frac{\pi(x-m)}{w}\right)\right) & \text{if } |x-m| < w \\ 0 & \text{otherwise.} \end{cases}$$

Again, we will use the following abbreviations:

$$\tau_i(x) := \tau_{b,y_i,a_i}(x) \qquad\qquad \tau'_i(x) := \tau_{b,y_i,c_i}(x)$$

Raised cosines add up equally nice as the triangular functions, which is illustrated in Figure 11. Thus, with the same parameters as before, we have $\tau_i(x) + \tau'_i(x) = y_i$ for $0 \le i < n$ and $x \in [a_i, c_i]$.

Example 5.3. *For our running example, we get the network in Figure 12 with each*  *unit replaced by a*  *unit using the same parameters.*

Similar to Theorem 5.2, one easily obtains the following result.

Theorem 5.4.

$$\sum_{i=0}^{n-1} \left(\tau_i(x) + \tau'_i(x)\right) = g(x) \quad \text{for all } x \in D.$$

5.3 Refining Networks

Our radial basis function network architecture lends itself to an incremental handling of the desired error bound. Assume we have already constructed a network approximating f_P up to a certain ε. We now want to increase the precision by choosing ε' with $\varepsilon > \varepsilon' > 0$, or by increasing the greatest relevant output level. Obviously we have $o_{\varepsilon'} \ge o_\varepsilon$ for $\varepsilon > \varepsilon' > 0$.

For this subsection, we have to go back to the original functions and domains from Section 3. Defining

$$\Delta P_{l_1,l_2} := \left\{ A \leftarrow \text{body} \in \mathcal{G}(P) \middle| l_1 < \|A\| \le l_2 \right\},$$

one can easily obtain the following result.

Lemma 5.5. If $l_2 \ge l_1$, then $\hat{l}_2 \ge \hat{l}_1$, $D_{l_2} \subseteq D_{l_1}$, $P_{l_2} = P_{l_1} \cup \Delta P_{l_1,l_2}$, and $P_{l_1} \cap \Delta P_{l_1,l_2} = \emptyset$.

Thus, the constant pieces we had before may become divided into smaller pieces (if the greatest relevant input level increases) and may also be raised (if any of the new clauses applies to interpretations represented in the range of that particular piece).

Looking at the body atoms in $\Delta P_{l_1,l_2}$, we can identify the pieces which are raised, and then add units to the existing network which take care just of those pieces. Due to the local receptiveness of RBF units and the properties of D_l stated above, the new units will not disturb the results for other pieces. Especially in cases where $|\Delta P_{l_1,l_2}| \ll |P_{l_1}|$, this method may be more efficient than creating a whole new network from scratch.

Example 5.6. *Assume that we want to refine the network in Figure 12 so that it approximates f_P up to a maximum error of* 0.007. *We obtain* $o_{0.007} = 4$, *and since our network is tailored for* $o_{0.02} = 3$, *we have to compute*

$$\Delta P_{3,4} = \{e(s^3(0)) \leftarrow \neg e(s^2(0)).\}$$

The clause in $\Delta P_{3,4}$ *has an effect on* $T_{P_4}(I)$ *if and only if* $e(s^2(0)) \notin I$. *In terms of the embedding, this means that* $f_{P_4}(x)$ *is different from (and larger than)* $f_{P_3}(x)$ *if and only if the* 3*rd digit of* x *is* 0, *since* $\|e(s^2(0))\| = 3$. *This implies that the domain of* $\hat{f}_{P_4}$ *is more fine-grained than the domain of* $\hat{f}_{P_3}$, *which is also reflected by the fact that* $\hat{4} = 3$ *and* $\lambda_4 = 0.0001111\ldots_3 = \frac{1}{54}$. *Put differently, the constant pieces of* $\hat{f}_{P_3}$ *are split up and some parts are raised. So now all we have to do is to determine which parts are raised and to add the corresponding units to the existing network.*

We consider the first 3 *digits of our inputs and are interested in those ranges where the* 3*rd digit is* 0. *Obviously, there are four such ranges, namely:*

$$D_{4,0} = [0.000000\ldots_3; 0.000111\ldots_3] \qquad D_{4,4} = [0.100000\ldots_3; 0.100111\ldots_3]$$
$$D_{4,2} = [0.010000\ldots_3; 0.010111\ldots_3] \qquad D_{4,6} = [0.110000\ldots_3; 0.110111\ldots_3]$$

Exactly on these ranges, $\hat{f}_{P_4}(x)$ *adds* $R\left(e(s^3(0))\right) = 0.0001_3$ *to the value of* $\hat{f}_{P_3}(x)$, *thus raising the four corresponding constant pieces by that value. So we have to add four pairs of units to our existing network, where on each of the above ranges exactly one pair computes the value* 0.0001_3. *We will omit the detailed computation of the parameters, which is completely analogous to the previous computations. Instead, Figure 13 just shows the graph of the function computed by the extended network.*

We could also right away construct the network for P_l by starting with one for P_1 and refining it iteratively using $\Delta P_{1,2}, \Delta P_{2,3}, \ldots, \Delta P_{l-1,l}$, or maybe using difference programs defined in another way, e.g. by their greatest relevant input level. This may lead to more homogeneous constructions than the method used in the previous subsections.

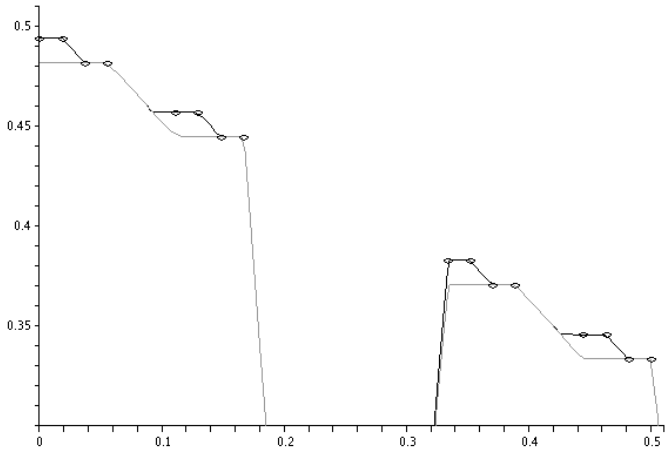

Figure 13: The graph of the function computed by the raised-cosine version of the original network from Figure 12 is shown in grey; the black parts are the results of the extensions described in Example 5.6. Circles denote the endpoints of the constant pieces of $\hat{f}_{P_4}$.

6 Conclusions and Future Work

In this paper, we have shown how to construct connectionist systems which approximate covered first-order logic programs up to arbitrarily small errors. We have thus, for a large class of logic programs, provided constructive versions of previous non-constructive existence proofs and extended previous constructive results for propositional logic programs to the first-order case.

An obvious alternative to our approach lies in computing the (propositional) ground instances of clauses of P up to a certain level and then using existing propositional constructions as in [11]. This approach was taken e.g. in [15], resulting in networks with increasingly large input and output layers. We avoided this for three reasons. Firstly, we want to obtain differentiable, standard architecture connectionist systems suitable for established learning algorithms. Secondly, we want to stay as close as possible to the first-order semantics in order to facilitate refinement and with the hope that this will make it possible to extract a logic program from a connectionist system. Thirdly, we consider it more natural to increase the number of nodes in the hidden layer for achieving higher accuracy, rather than to enlarge the input and output layers.

In order to implement our construction on a real computer, we are facing the problem that the hardware floating point precision is very limited, so we can only represent a small number of atoms in a machine floating point number. If we do not want to resort to programming languages emulating arbitrary precision, we could try to distribute the representation of interpretations on several units, i.e. to create a connectionist system with multi-dimensional input and output. For real applications, it would also be useful to further examine the possibilities for incremental refinement as illustrated in Section 5.3.

Another problem is that the derivative of the raised-cosine function is exactly 0 outside a certain range around the peak, which is not useful for training algorithms like backpropagation. Gaussian activation functions would be more suitable, but appear to be much more difficult to handle.

We are currently implementing the transformation algorithms, and will report on corresponding experiments on a different occasion. One of our long-term goals follows the path laid out in [7, 5] for the propositional case: to use logic programs as declarative descriptions for initialising connectionist systems, which can then be trained more quickly than randomly initialised ones, and then to understand the optimised networks by reading them back into logic programs.

References

[1] Krzysztof R. Apt, Howard A. Blair, and Adrian Walker. Towards a theory of declarative knowledge. In Jack Minker, editor, *Foundations of Deductive Databases and Logic Programming*, pages 89–148. Morgan Kaufmann, Los Altos, CA, 1988.

[2] Sebastian Bader, Artur S. d'Avila Garcez, and Pascal Hitzler. Computing first-order logic programs by fibring artificial neural networks. In *Proceedings of the 18th International FLAIRS Conference, Clearwater Beach, Florida, May 2005*, 2005. To appear.

[3] Sebastian Bader and Pascal Hitzler. Logic programs, iterated function systems, and recurrent radial basis function networks. *Journal of Applied Logic*, 2(3):273–300, 2004.

[4] Sebastian Bader, Pascal Hitzler, and Steffen Hölldobler. The integration of connectionism and knowledge representation and reasoning as a challenge for artificial intelligence. In L. Li and K.K. Yen, editors, *Proceedings of the Third International Conference on Information, Tokyo, Japan*, pages 22–33. International Information Institute, 2004. ISBN 4-901329-02-2.

[5] Artur S. d'Avila Garcez, Krysia Broda, and Dov M. Gabbay. Symbolic knowledge extraction from trained neural networks: A sound approach. *Artificial Intelligence*, 125:155–207, 2001.

[6] Artur S. d'Avila Garcez, Krysia B. Broda, and Dov M. Gabbay. *Neural-Symbolic Learning Systems — Foundations and Applications*. Perspectives in Neural Computing. Springer, Berlin, 2002.

[7] Artur S. d'Avila Garcez and Gerson Zaverucha. The connectionist inductive lerarning and logic programming system. *Applied Intelligence, Special Issue on Neural networks and Structured Knowledge*, 11(1):59–77, 1999.

[8] Michael Gelfond and Vladimir Lifschitz. The stable model semantics for logic programming. In Robert A. Kowalski and Kenneth A. Bowen, editors, *Logic Programming. Proceedings of the 5th International Conference and Symposium on Logic Programming*, pages 1070–1080. MIT Press, 1988.

[9] Pascal Hitzler, Steffen Hölldobler, and Anthony K. Seda. Logic programs and connectionist networks. *Journal of Applied Logic*, 2(3):245–272, 2004.

[10] Pascal Hitzler and Anthony K. Seda. Generalized metrics and uniquely determined logic programs. *Theoretical Computer Science*, 305(1–3):187–219, 2003.

[11] Steffen Hölldobler and Yvonne Kalinke. Towards a massively parallel computational model for logic programming. In *Proceedings ECAI94 Workshop on Combining Symbolic and Connectionist Processing*, pages 68–77. ECCAI, 1994.

[12] Steffen Hölldobler, Yvonne Kalinke, and Hans-Peter Störr. Approximating the semantics of logic programs by recurrent neural networks. *Applied Intelligence*, 11:45–58, 1999.

[13] John W. Lloyd. *Foundations of Logic Programming*. Springer, Berlin, 1988.

[14] R. Rojas. *Neural Networks — A Systematic Introduction*. Springer, 1996.

[15] Anthony K. Seda and Máire Lane. On approximation in the integration of connectionist and logic-based systems. In L. Li and K.K. Yen, editors, *Proceedings of the Third International Conference on Information, Tokyo, Japan*, pages 297–300. International Information Institute, 2004. ISBN 4-901329-02-2.