Results from Analogy Experiment

Means and variances of dot-produc

Chapter Since the early 1980's there has been considerabl

higher-level cogni8ive tasks such as natur1l language pro

if we want to understand the difficulties and possibiliti

Much of the interest in connectionist models for hi

tion, approxim18e matching, good gener1lization, gracef

isthatmanyhigher-levelreasoningproblemsappeartoa

I assume that the reader is familiar with the basic p

Representations, descriptions, and implementations

In this thesis I am more concerned with how complex data structures can be elegantly in distributed representations than with the choices i

models. The first is in the activation values of the units. are used for representing short-term data; the particular instance of the task, ty

message passed on to other units. nearly all connectionist models, units can (in compute their values in parallel, and the overall computation time is short – on the or of 100 steps or less.

There are two distinct locations for representing knowledge (or more neutrally, data) in

One is as constraints on solutions to a task, the other is as transformations between and output are often used with those localist models that process by

settling to a stable or low-energy state. These states are (or are hoped to be) those in Constraints most constraints are In the network the constraints represented by a link are either the connected units should

Pattern transformations are generally used in feedforward and recurrent and are usually learned from examples. Some networks, such the Hopfield [1982] network

(an error-correcting associative memory) can be interpreted in both ways.

Theris far more diversity of representation schemes for the temporary data structur

suitability for learning are central to this There are other important but less prissues which I mention here but do not return **Represe**

in the links and determines what the network can Representational compute.Representations are of critic1l importance in connectionist To a large extent, representation scheme used

arthe source of both strengths and weaknesses of connectionist models. Carstructed connectionist representations can endow systems with useful properties s

prtend to be more associ18ed with distributed representations than with representations (though loc1l representations are not without How

appropri18e interpretation of vowel sound can depend on the preceding amount of input that must b

Connectionist representations are often very different from convention1l data tures,such as v

unintelligible and appear to have little or no relationship to convention1l data structurHowever, if the mo

perform some task, one must consider what sort of structures the model is going to reprand how it will manipulate them. designing or analyzing reprfor complex structures, a number of questions concerning several of representation gen

about how to perform the

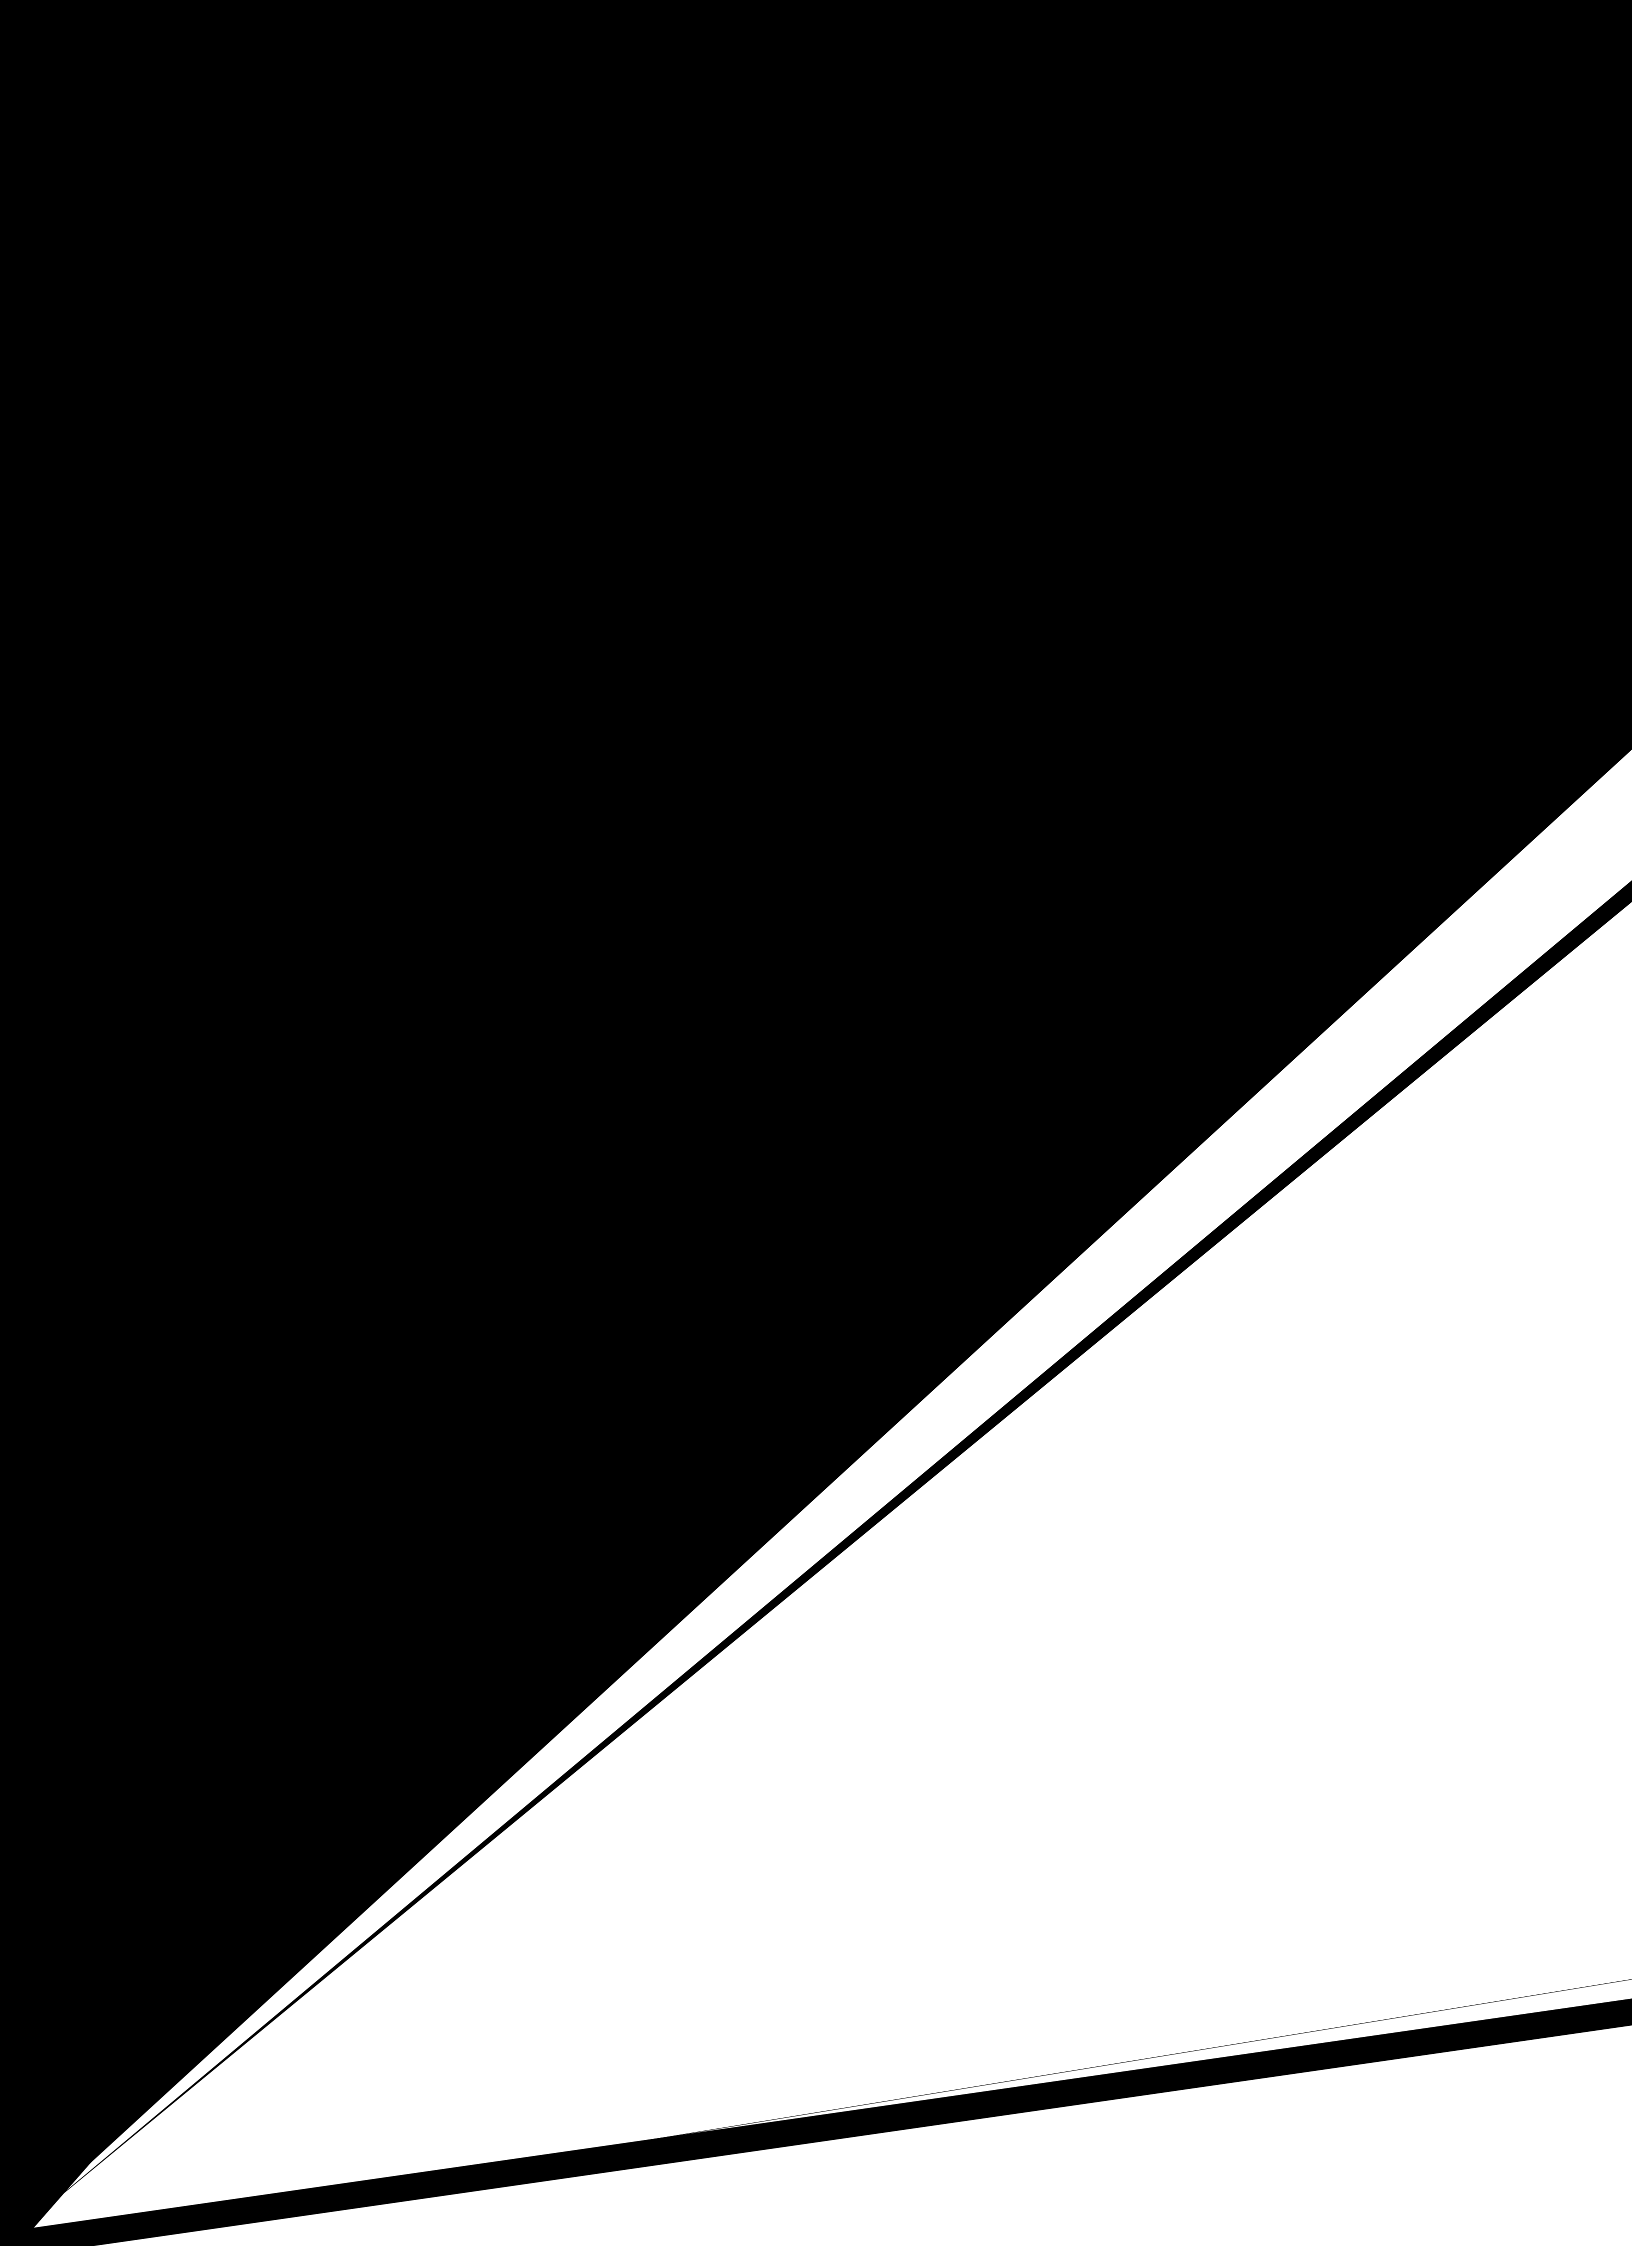

is to follow the pointerHash indexing is a technique from convention1l data-structurtheory which

providemoreexplicitlabelsforconcepts. could make it possible to process morefsince it would not always be necessary to follow pointers to information abou

themworthconsideringasarepresentationforhigher-levelcognitivetasks. In Chapter 2, Pollackreview various connectionist models that are relevant to the representation a

In Chapter 3, I present and discuss a representation1l scheme for hierarchic1l structurversion.called “Holographic Reduced Representations” (HRRs), wh

useful forms of similarity. is obvious from first principles that HRRs will some sort of similarity structure on structured objects, since

address bits are used to indic18e the type of the object at that address, e.g., atom or cons This helps increase the speed of LISP pr3

MacLennan [1991] claims this is impossible to do with continuous representations, due to a theorem Brouwer’s which sta

$\times$ n $\rightarrow$ n with a continuous $-: \rightarrow \times$.

Forbus's [1989] notions of the analogical similarity of structured object

Chapter Review of connectionist distributed memory mo

In this chapter, I review various connectionist models

Loc1list networks that compute by spreading activation an
the first connectionist models of higher-level A unit in o

represent18ions are gener1lly loc1l in the sense that units represent propositions in on

sometimes loc1l representations can be considered to be at a higher level, at which the units are the features of Derthick [1990] this point, and his

which requires more units. The downside of this simple and explicit style
of representation is that networks can become very large when there are many propositions. Holyoak and Thagard's [1989] A

is an example of this; the propositions are hypothesis about matches between structure elements, and there are many

constructed by hand or by a computer program from a priori knowledge, rather than learned from Shastri's system is a

because it requires complex tempor1ry data structures, and more processing power simple provided by a finite state

is to demonstrate a system that could apply all of these constraints in par1llel during the consider this style of processing more efficient and more cally plausible than rule-based sys

by semantic an1lysis, psychologi-Their model uses different units to represent different syntactic

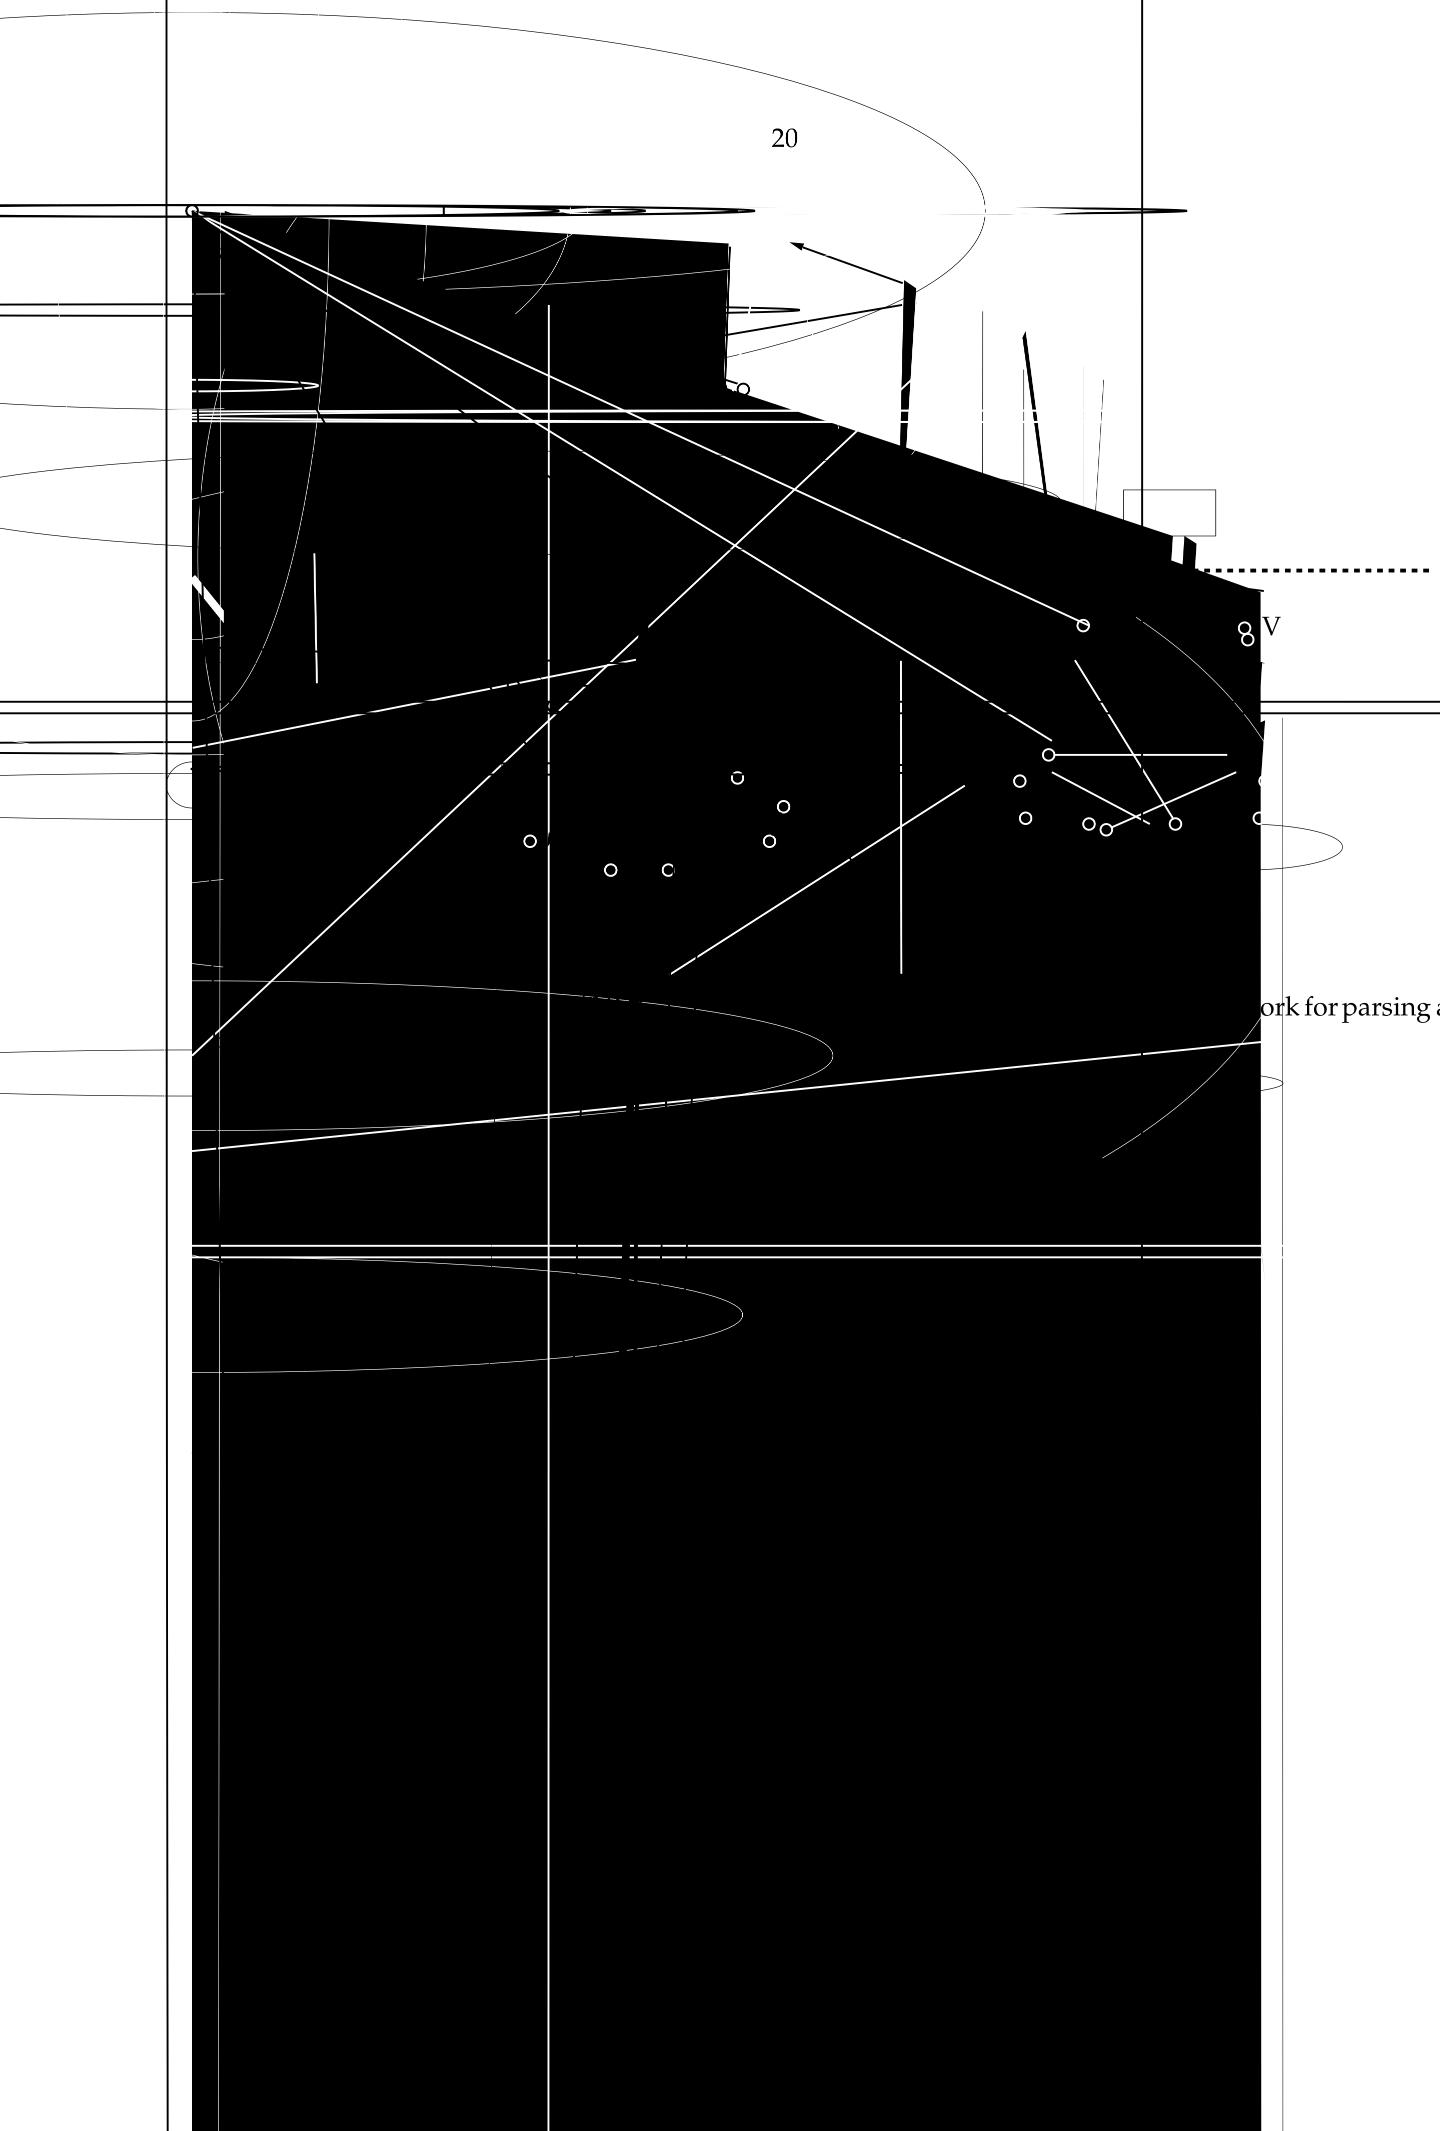

ork for parsing and interpreting

(b)

S

S

VP

S → NP Figure

primitive

0

VP → verb PP

unit

2

2

2

1

1rule

VP

shorthand notation

verb

8 9

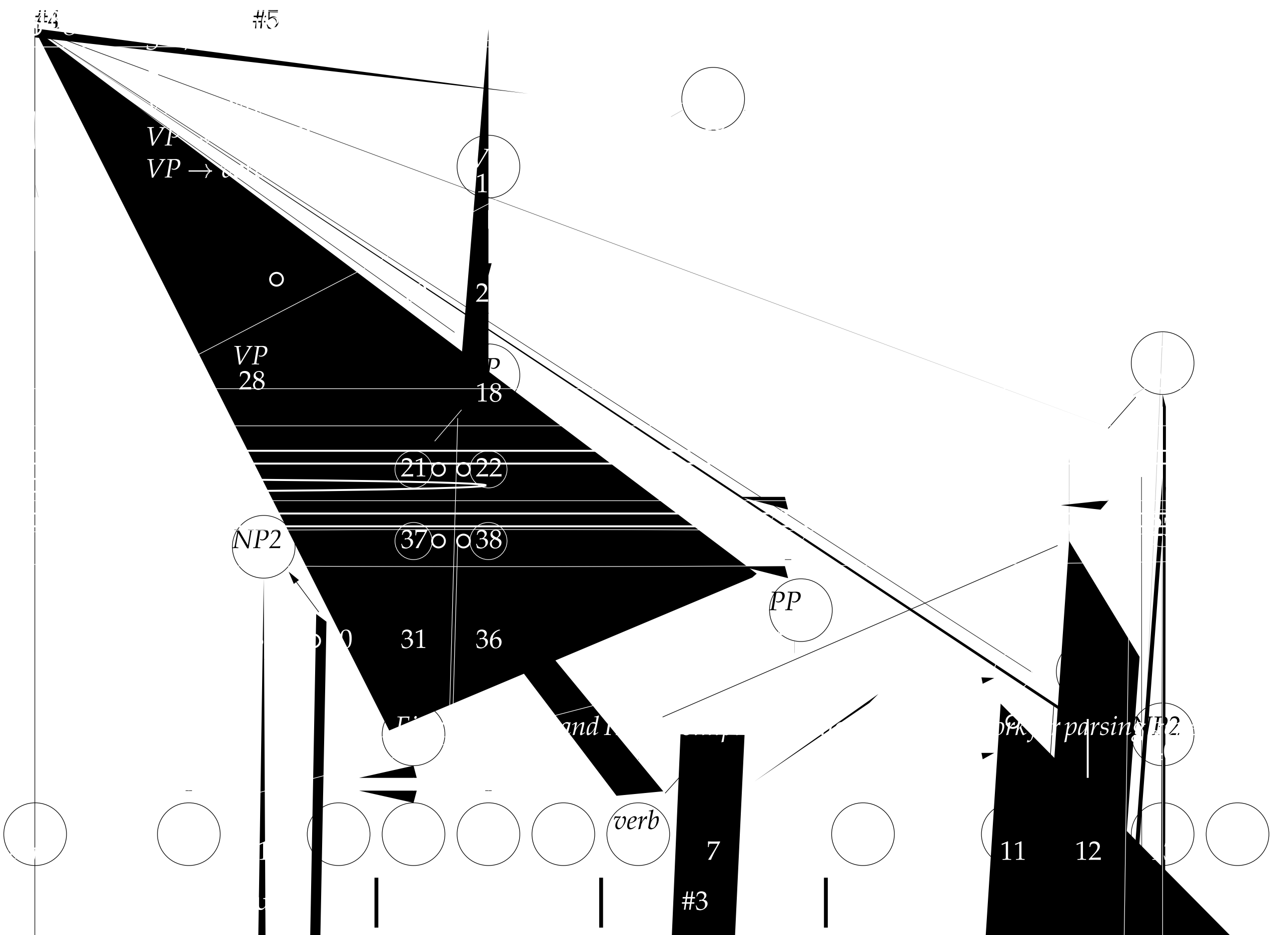

Fi and I ork for parsing

1 5]

T e hes th rrect

ex c heds. n use very simple computation because it has

he t n of th

ar T minal can appear in the parse tree. For example, Node3

h ra

ob n a hard th of 2.3:it can parse. Another is that the numbe

time to grows with

afru at ibrium 2 ere are high energy barriers, Boltzmann machines take a tha

In pra e run for a finite

n a ch to guaranteed if the temperature 42time k, $T(k)$, satisf

roport n height of energy barriers. we are interested in the