

Dynamics of learning: A case study


Jay McClelland

Stanford University

Piaget's theory of stages and development

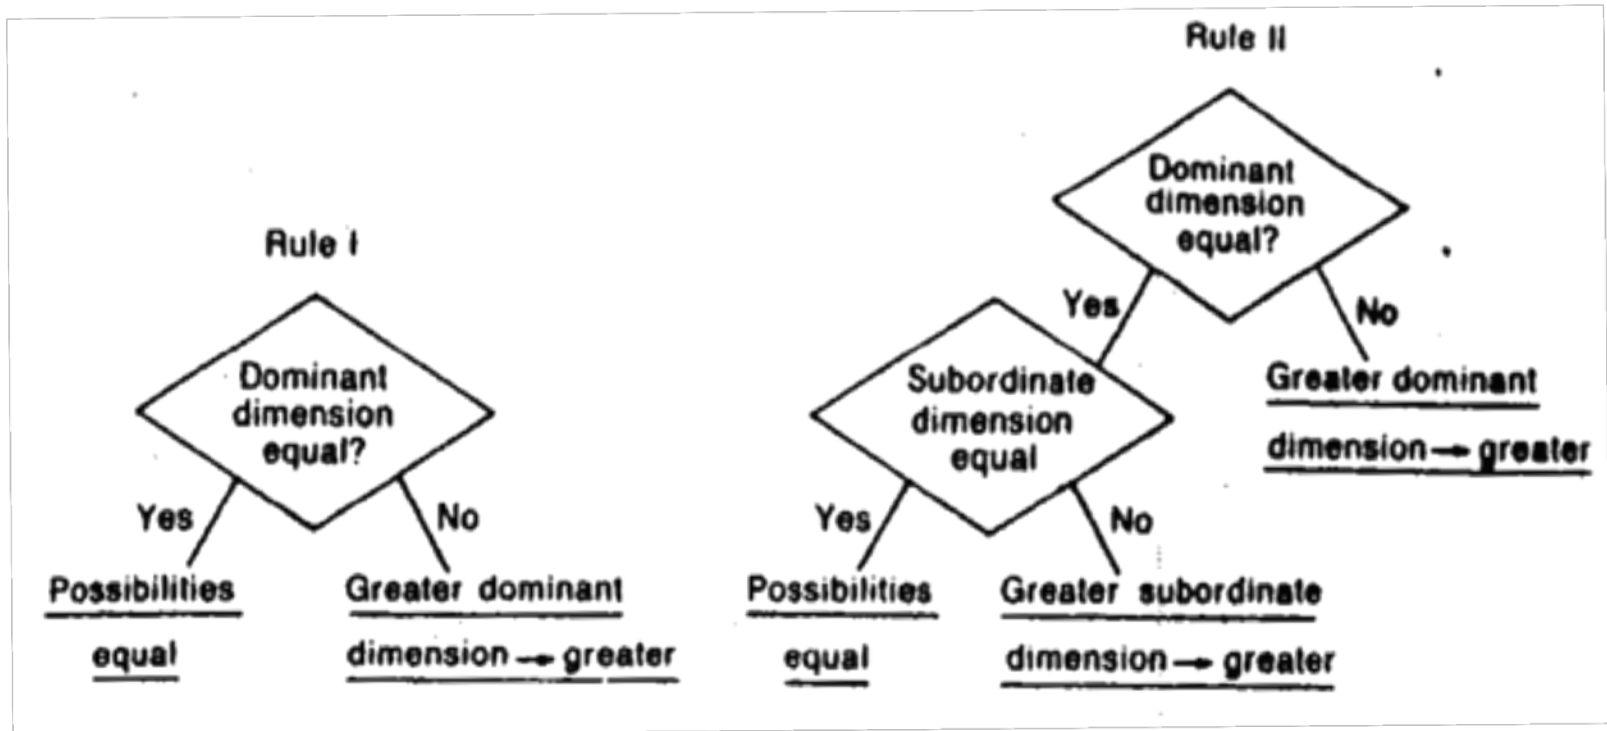
- Our knowledge takes qualitatively different forms at different stages
- This occurs because our knowledge schemas adapt with experience
 - ‘Equilibration: Assimilation & Accommodation’
- Questions about the theory:
 - What do these words mean?
 - How does it work?
 - How can incremental adjustments lead to qualitative change?

An example domain: The balance scale



Will the scale tip left, tip right, or balance?

Siegler's First Two Rules



Rule assessment procedure:

24 problems of 6 types


Children are characterized as 'using a rule'
if 20/24 responses are consistent with it

Findings

- Children 5-7 generally behave in accordance with Rule 1
- Then, most children's behavior is briefly consistent with Rule 2, before they progress to more complex patterns of responding
- Around the time of these transitions, children exhibit increasing, graded sensitivity to the distance cue
 - Rule 1 kids will take distance into account if the distance difference is large enough

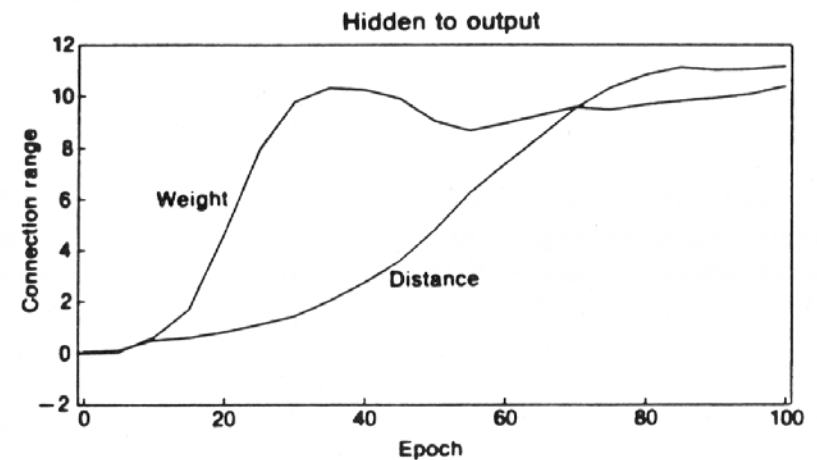

Siegler's Intervention Experiment

- Two groups of 'Rule 1' children
 - Young (about 5 yrs old)
 - Older (7-8 yrs old)
- All children see 16 'conflict' problems
- After the experience:
 - Older children progress to Rule 2
 - Younger children show no change or start choosing randomly




An Alternative to Rules

- Knowledge is implicit
 - In the connections in a neural network
- Experience with problems shapes the connections
- The network learns to become sensitive to weight and distance gradually through the accumulation of small changes
- Like children, it shows graded sensitivity to distance



Graded Differences in Connections Lead to Different Outcomes of Learning at Different Points in Development

- Training with conflict problems at epoch 20 leads to degradation of performance.
- Training at 40 epochs leads to progress.



Is the Model too Continuous?


- ‘Rules’ might be seen as similar to phases of matter (solid, liquid, gas)
- This view predicts that children’s performance should exhibit ‘catastrophe flags’ that signify such transitions


Another Interesting Question


Can the model teach
itself?


Van der Maas *et al* Experiment


- Pre-test
- A series of distance problems
- Post-test
- Over 300 children tested!
























Van der Maas *et al* Catastrophe Flags

- Pre and post-test
 - ‘Inaccessible region’
 - Jump across gap


- Distance series
 - ‘Maxwell’
 - ‘Hysteresis’
 - ‘Sudden jump’

Findings

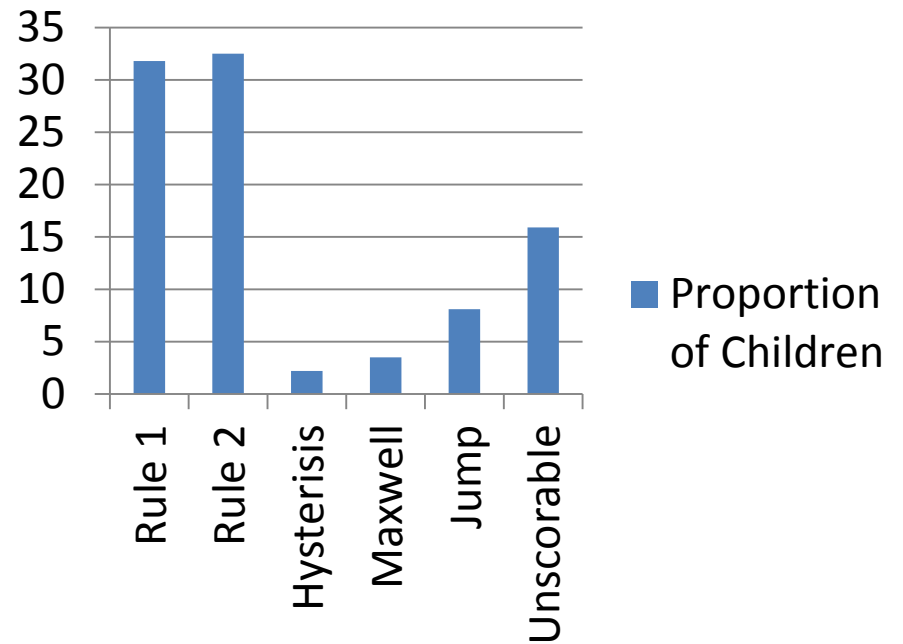
Pretest Scores



Posttest Scores




Distance Series: Proportion of Children Showing Each Type of Pattern




Simulation with Model that Teaches Itself

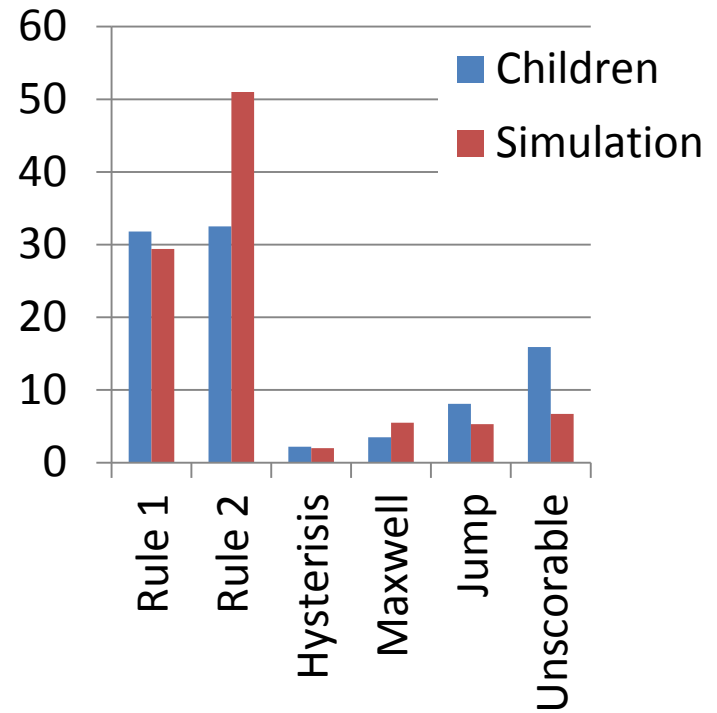
Model Pretest Scores



Model Posttest Scores



Distance Series: Children vs. Model



Conclusion

- We are far from a perfect understanding of transitions in learning
- But we have come a long way
- Simple rule-based models miss a lot
- Models with learning dynamics capture quite a lot more:
 - They show gradual knowledge accumulation and yet capture patterns of change similar to what we see in children