

The HUME Model-Driven Discovery System

Adrian Gordon¹

Laboratoire de Recherche en Informatique
Université de Paris-Sud
Orsay
France

SMTP% "adrian.gordon@unn.ac.uk"

Abstract

Structural models provide an important source of hypothetical knowledge in scientific discovery. Informal Qualitative Models (IQMs) are structural models which can be applied to weak theory scientific domains. Example models are presented for the domain of solution chemistry. These models can be systematically generated, but, due to the weak theory nature of the domains to which they are applied, they cannot be verified directly. Instead, the application of IQMs to a problem can be used to drive other scientific discovery processes; in particular, the discovery of numeric laws. The HUME system is a discovery system based around the application of IQMs. HUME's discovery goal is to construct explanations for phenomena, such as the depression of the freezing point of salt solutions, using a variety of reasoning strategies. HUME first attempts to explain such phenomena using a pre-existing theory. If this theory is not able to provide an explanation, the system uses a combination of theory construction and numeric law discovery. The application of IQMs provides hypotheses for use by the other two processes. Used in this way, IQM application can be seen to provide a degree of explanatory support for numeric laws which would otherwise be simply descriptive generalisations of data. An example of the application of HUME to a problem in solution chemistry is presented.

Informal Qualitative Models in Scientific Discovery

Structural models can provide an important source of hypothetical knowledge for use in scientific discovery. *Informal Qualitative Models (IQMs)* are one sort of structural models that are used by scientists. IQMs are abstract structural descriptions of physical systems, either actual or hypothetical. They were first introduced in Sleeman *et al.* (1989), and further elaborated in Gordon (1992, 1993) and Gordon *et al.* (1994, 1995).

Figure 1 gives an example of two models which can be used in the domain of solution chemistry. Model A shows the case where molecules of solvent and solute are uniformly physically distributed throughout a solution,

with neither sets of molecules being chemically changed in any way. This is the *Physical Mixing* model. Model B shows the *Association* model, in which molecules of solute form an *association* with molecules of solvent. The associations thus formed are dissolved in the remainder of the solvent. An example might be a salt which existed in a *hydrated* form in solution.




Figure 1. The (A) Physical Mixing and (B) Association IQMs

A historical study of eighteenth and nineteenth century solution chemistry has led to the elaboration of a set of increasingly complex IQMs. As we show in (Gordon *et al.* 1994) each of these models can be generated from the simplest model in the domain, the *Physical Mixing* model, by the application of a set of *model generation operators*. The repeated application of these operators results in the synthesis of a *search space* of IQMs.

Clearly, this formulation of a search space of structural models is closely related to previous work in computational discovery. Systems such as STAHL/DALTON (Langley *et al.* 1987), REVOLVER (Rose and Langley, 1986) and BR3 (Kocabas, 1991) all formulate discovery as a heuristic search through a space of models.

However, previous approaches have used strong heuristics to constrain the generation of models, and to confirm their validity. BR3, for example, uses a set of *quantum conservation laws* to constrain the generation of quark models in particle physics. Many previous systems are also able to *confirm the observational consequences* of generated models. This is frequently done by looking for observed reactions (such as particle decay reactions or chemical reactions) which confirm the generated models. In weak theory domains, these constraints do not operate. None of the models used historically in solution chemistry had directly observable consequences, for example. In early solution chemistry history, IQMs were confirmed by looking for numeric laws. If a "good" numeric law could be found based on the applied model, then this was accepted as evidence for the validity of the model.

The HUME Discovery System

HUME is a discovery system which is based around the application of IQMs. The view of scientific understanding

¹ Current Address: Department of Computing, University of Northumbria at Newcastle, Newcastle-Upon-Tyne, NE1 8ST, UK.

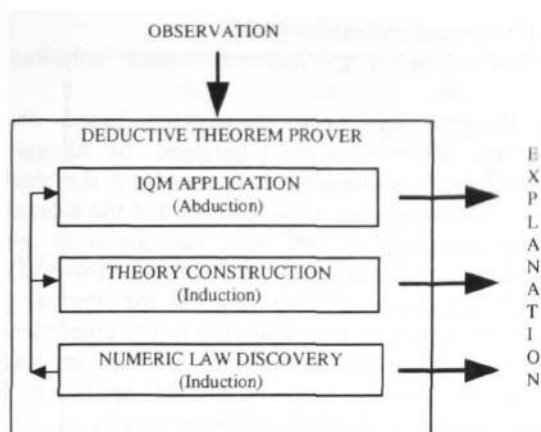


Figure 2. General Architecture of HUME

that is adopted by HUME is that it depends on the construction of *deductive nomological explanations* (Hempel and Oppenheim, 1948). HUME attempts to understand a phenomenon by explaining it in terms of an existing domain theory. However, since HUME will be in most cases dealing with *incomplete* domain theories, a significant sub-task for the system is the construction or extension of domain theories.

Figure 2 illustrates the general architecture of HUME. The system is based around a deductive theorem prover. This theorem prover takes an observation as input, and produces an explanation for this observation on successful exit. The explanation is based on the application of a domain theory, and background domain knowledge. However, if this domain theory cannot initially provide an explanation, the other modules of HUME are called to extend that domain theory. These modules are an IQM application module, which employs basically *abductive* reasoning techniques; a theory construction module, based around that of the DISCIPLINE system (Tecuci and Kodratoff, 1990); and a numeric law discovery module, which is at present provided by the ARC² system (Moulet, 1991). The discovery element of the system lies not merely in the construction of explanations, but in the *confirmation of existing IQMs, the construction of new domain theories, and the discovery of new laws.*

Table 1. Freezing point data. Asterisks indicate *anomalous datapoints* (See later explanation).

| Substance | Formula | Molecular Mass (M) | Freezing point °C (Fp) | M × Fp |
|--------------------|--|--------------------|------------------------|--------|
| Barium Chlorate | Ba(ClO ₃) ₂ | 304 | -0.145 | -44.1 |
| Barium Nitrate | Ba(NO ₃) ₂ | 261 | -0.155 | -40.5 |
| Strontium Nitrate | Sr(NO ₃) ₂ | 211 | -0.195 | -41.2 |
| Copper Nitrate | Cu(NO ₃) ₂ | 187.2 | -0.244 | -45.7 |
| Copper Acetate* | Cu(C ₂ H ₃ O ₂) ₂ | 181 | -0.171 | -31.1 |
| Lead Acetate* | Pb(C ₂ H ₃ O ₂) ₂ | 325 | -0.068 | -22.2 |
| Barium Chloride | BaCl ₂ | 208 | -0.233 | -48.6 |
| Mercuric Chloride* | Hg Cl ₂ | 271 | -0.076 | -20.5 |

²ARC is an extension of the ABACUS system

HUME's inputs are a set of experimental observations, a set of theorems which constitutes a background domain theory, and a library of IQMs. All of this knowledge is represented in Horn clause form. A discovery problem must then be specified to the system. This discovery problem takes the form of a particular experimental problem, and the result of that experimental problem. Typically, this will be *one* of the experimental observations that the system has already seen.

The direct output from HUME on the successful completion of a discovery task is an *explanation* for a specified discovery problem, in the form of a proof tree. However, though a proof tree is the *direct* result of successful discovery, there may be many *indirect* results. Some of the theorems used in the proof tree might themselves have been constructed by HUME, and some of the "facts" used in the proof tree might be hypotheses introduced by the system.

An Example: HUME and Solution Chemistry

This section will describe how HUME can be used to undertake a discovery task in the domain of solution chemistry. The data used in this example discovery session are taken from (Raoult, 1885, pp. 407-408). This data deals with the freezing points of a set of thirty aqueous solutions of the salts of various bivalent metals. Table 1 shows some of this data. Table 1 shows the name and empirical formula for some of the salts used by Raoult, together with their molecular masses (M) and the observed freezing point (Fp) for a solution of each salt. In each case the solution involves 1g of solute dissolved in 100g of solvent. The calculated term $M \times Fp$ is also shown. HUME's initial discovery goal is to explain the observed freezing point of the first example from Table 1, a solution of Barium Chlorate, with a freezing point of -0.145 °C.

Initially, HUME is provided with knowledge about the *solvents* and *solutes* used in each example from Table 1, such as their composition, mass, and the molecular mass of the substances involved. However, the only knowledge provided for each resulting *solution* is its observed freezing point. HUME's initial domain theory is therefore incomplete, and the system is initially unable to explain the observed freezing point of the selected example. The system's discovery modules are called to extend the initial theory. The system first attempts to use its theory construction strategy, adapted from that of the DISCIPLINE system. The basis of this strategy is a search for possible *causal relationships* amongst the objects implicated in a problem. This process is described in detail in (Tecuci and Kodratoff, 1990) and (Gordon, *in preparation*). However, since the system initially has knowledge only about the composition of the *solvent* and *solute* used in the first example, and not the resulting *solution*, it is unable to discover

(Falkenhainer and Michalski, 1987).

any possible causal relationships. At this stage, the system can proceed no further without making some *hypotheses* about the structure of the solution. This is done by applying an IQM.

Each IQM in HUME is represented in the form of a Horn clause theorem. Theorem 1 represents the *Physical Mixing* IQM, for example.

Theorem 1:

```
(← (model physical-mixing ?eg)
  (and (make-solution ?eg ?solvent ?solute ?solution)
    (composed-of ?solvent [?substance1])
    (composed-of ?solute [?substance2])
    (composed-of ?solution [?c1 ?c2])
    (composed-of ?c1 [?substance1])
    (composed-of ?c2 [?substance2])))
```

where the symbol “←” represents logical implication, square brackets represent lists (Prolog syntax) and “?” preceding a symbol represents a variable. This theorem states that the *Physical Mixing* IQM applies to an example of the creation of a solution (“make-solution”), where the solvent is composed of some substance, ?substance1, the solute is composed of some substance, ?substance2, and the resulting solution is composed of two sub components (?c1 and ?c2) which are themselves composed of ?substance1 and ?substance2 respectively. In order to apply this IQM to a solution, the theorem is first matched against HUME’s knowledge base. Typically, this will result in a partial match, where some of the assertions required to fully instantiate the antecedents of the theorem are missing from HUME’s knowledge base. HUME’s *abductive* reasoning strategy in these circumstances simply asserts any missing antecedents directly into the knowledge base as *hypotheses*. In the current case, since no knowledge is initially available concerning the components of the resulting Barium Chlorate solution, the final three antecedents of the theorem will fail to match assertions in HUME’s knowledge base. These three antecedents will then be directly asserted into HUME’s knowledge base, with system generated symbols to replace any still uninstantiated variables. This amounts to hypothesising the unknown components of the solution.

Once an IQM has been applied in this way, HUME’s theory construction strategy can be applied again. This time, however, the system can make use of the hypothetical knowledge about the solution that has been introduced by the application of the *Physical Mixing* IQM. Theorem 2 results:

Theorem 2:

```
(← (freezing-point ?solution ?x)
  (and
    (composed-of ?solution [?c1 ?c2])
    (composed-of ?c1 [?s1])
    (composed-of ?solvent [?s1])
    (composed-of ?c2 [?s2]))
```

```
(composed-of ?solute [?s2])
(make-solution ?g1 ?solvent ?solute ?solution)))
```

This theorem has been constructed based on the relationships which now exist between the solvent, the solute, and the hypothesised components of the resulting solution. That is, it is based on the fact that the solvent and solute are composed of the same substances as are the hypothesised components of the solution. However, this theorem is not complete. The value of the freezing point property, ?x, is nowhere instantiated in the antecedents of the theorem. In circumstances such as this, in order to instantiate the missing property, HUME applies its third discovery strategy: the search for numeric laws.

In order to search for numeric laws, HUME uses partial theorems such as Theorem 2 as *contexts* for numeric law discovery. All of the numeric properties of objects which are implicated in the partial theorem are gathered together and passed to HUME’s law discovery module. However, this is not done only for the example currently under investigation, but for *all* examples known to the system. All of the examples from Table 1 would be used in this case, for example. Using this data, HUME’s numeric law

discovery module discovers the law $Fp = \frac{-42.986}{M}$ where

Fp is the freezing point of the solution, and **M** is the molecular mass of the solute³. Figure 3 shows the graph of freezing point against molecular mass for Raoult’s data. Figure 3 also shows the newly discovered law superimposed over this data. Once discovered, this law is then incorporated into the previously seen partial theorem, to produce Theorem 3:

Theorem 3

```
(← (freezing-point ?solution (/ -42.986 ?m)
  (and
    . . . antecedents as in Theorem 2
    . . .
    (molecular-mass ?s1 ?m)))
```

Once this theorem is constructed, HUME is able to construct an explanation for the observed freezing point of the first solution from Table 1. The resulting proof tree is shown in Figure 4. The basis of this proof tree lies in the application of Theorem 3, but some of the grounded assertions in this proof tree are only available because of the abductive application of the *Physical Mixing* IQM. These are the assertions shown underlined in Figure 4. These assertions concern the hypothetical composition of the solution (named *solution1* in this case).

Conclusions

The previous section showed how HUME was able to undertake a discovery task in the domain of Solution Chemistry. The system was able to explain the observed

³This law can also be expressed as $Fp \times M = -42.986$. Values for $Fp \times M$ are shown in Table 3.




Figure 3. Graph of Freezing point against molecular mass for Raoult's data.

freezing point of a particular example solution. However, the theorem that the system constructed, and the numeric law embodied in that theorem, are also able to explain the observed freezing points of a significant number of other solutions from Table 1. Thus, by starting from the careful analysis of a single example the system is able to construct a generally useful domain theory for a particular aspect of solution chemistry.

Although the system uses a number of different modes of reasoning, the general discovery strategy is guided by the application of IQMs. IQM application provides hypotheses for use in theory construction, and frequently allows the derivation of numeric data for use in law discovery (Gordon, 1992, 1993). Furthermore, IQMs also lend a degree of explanatory support for discovered numeric laws. In the case of Theorem 3 above, the two parts of the theorem are mutually supporting. The explanatory component represented by the antecedents of the theorem (which have been constructed after IQM application) supports the discovered numeric law. Similarly, the




Figure 5. The Solute Dimerism IQM

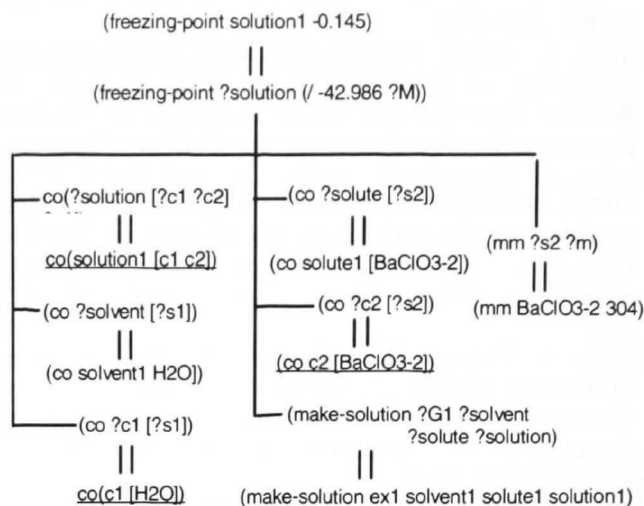


Figure 4. Proof tree for "(freezing-point solution1 -0.145)". mm = "molecular mass", co = "composed-of", BaClO3-2 = Barium Chlorate, H2O = water.

discovered numeric law also serves to confirm the validity of the applied IQM.

IQMs can also explain away anomalies in a domain theory. As Figure 3 shows, there are three examples in the data of Table 1 which appear to be anomalous with regard to the numeric law discovered by HUME. That is, the observed value of the freezing point of each of these examples differs significantly from that predicted by the law (these examples are labelled with asterisks in Table 1). As can be seen in Table 1 the three anomalous examples have a value for the term $Fp \times M$ which is roughly half of that observed for the other examples. However, if we assume that the molecules of each of the substances in the anomalous cases exist in solution associated two by two, then the value of their molecular masses, M , would be doubled in each case, and the substances would then fit the general law fairly closely. Hypothesising that these salts associate two by two in solution in fact corresponds exactly to the application of a new IQM to each of these examples. This is the Solute Dimerism model shown in Figure 5. Raoult himself explained away these apparent anomalies in precisely this way. A careful study of the history of solution chemistry offers numerous other examples in which new IQMs, or variations on existing IQMs were used to derive and justify numeric laws, and explain away seemingly anomalous observations (Gordon, *in preparation*).

References

- Falkenhainer, B.C. and Michalski, R. S. (1986). Integrating Quantitative and Qualitative Discovery: The ABACUS System. *Machine Learning*, 1, pp. 367-401.
- Gordon, A. (1992). Informal Qualitative Models and Scientific Discovery. In *Proceedings of the ML92*

- Workshop on Machine Discovery*, J. M. Zytkow (Ed), pp. 98-102, Aberdeen.
- Gordon, A. (1993). Informal Qualitative Models and the Depression of the Freezing Point of Solutions. In *Working Notes for the MLnet Workshop on Machine Discovery*, P. Edwards (Ed), pp. 56-60, Blanes, Spain.
- Gordon, A. Edwards, P., Sleeman, D. and Kodratoff, Y. (1994). Scientific Discovery in a Space of Structural Models: An Example from the History of Solution Chemistry. *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*. pp. 381-386.
- Gordon, A., Sleeman, D. and Edwards, P. (1995). Informal Qualitative Models: A Systematic Approach to their Generation. *Working Notes for the AAAI Spring Symposium on Systematic Methods of Scientific Discovery*. pp. 19-22.
- Gordon, A. *Informal Qualitative Models in Scientific Discovery*. Thèse de Docteur en Sciences, Université de Paris-Sud, Centre D'Orsay, *in preparation*.
- Hempel, C. and Oppenheim, P. (1948). Studies in the Logic of Explanation. *Philosophy of Science*, vol 15. pp. 135-175.
- Kocabas, S. (1991). Conflict Resolution as Discovery in Particle Physics. *Machine Learning*, 6, pp. 277-309.
- Langley, P., Simon, H.A., Bradshaw, G.L., and Zytkow, J.M. (1987) *Scientific Discovery: Computational Explorations of the Creative Processes*. Cambridge, MA: MIT Press.
- Moulet, M. (1991). Using accuracy in law discovery. In *Proceedings of the Fifth European Working Session on Learning*. Y. Kodratoff (Ed), Lecture Notes in Artificial Intelligence, Springer Verlag, pp. 118-136.
- Raoult, F.M. (1885). Sur le point de congélation des dissolutions salines. *Annales de Chimie et de Physique*, (6) IV, pp. 401-430.
- Rose, D., and Langley, P. (1986). Chemical Discovery as Belief Revision. *Machine Learning*, 1, pp. 423-452.
- Sleeman, D.H., Stacey, M.K., Edwards, P. and Gray, N.A.B. (1989). An Architecture for Theory-Driven Scientific Discovery. In *Proceedings of the Fourth European Working Session on Learning* (pp. 11-24). Montpellier, France: Pitman/Morgan-Kaufmann.
- Tecuci, G. and Kodratoff, Y. (1990). Apprenticeship Learning in Imperfect Domain Theories, in *Machine Learning: An Artificial Intelligence Approach*, Vol. 3, Y. Kodratoff, and R.S Michalski (Eds.), pp. 514-552, Morgan Kaufmann, San Mateo CA.