Automating Representation Change Across Domains for Reasoning*

Aaron Stockdill[0000−0003−3312−5267]

University of Cambridge, Cambridge, UK
aaron.stockdill@cl.cam.ac.uk

Introduction. Changing the way we represent the problem can have powerful benefits, both pedagogically and cognitively [1,2]. The common advice to 'draw a picture' prompts deep questions about the nature of representations and problem solving: what makes up a representation, what should be considered when choosing a representation, and how do we evaluate a representation? Our work explores each of these questions, motivating a framework for the discussion of representations both individually, and the links between them.

The main contribution of this work is the *correspondence* framework for describing the relationships between *components* of representational systems, and allowing us to compute the *informational suitability* and *cognitive cost* of a particular representational system with respect to both a specific problem being solved, and the particular person solving the problem. This work is part of the broader rep2rep project at the University of Cambridge (with Prof. Mateja Jamnik, my supervisor) and the University of Sussex (with Prof. Peter Cheng).

Representing representations. This research project introduces *components* and how they can be composed into *descriptions* of problems, representations, and representational systems [4]. Components are the building blocks of representations, defined as triples of *kind*, *value*, and *attributes*. The 'kind' partitions components into five categories, three notational—primitives, types, and patterns—and two inferential—laws and tactics. Attributes give extra information about components, linking them with other components through relationships such as 'is-of-type'. Representations may be described by many components, or very few components; components may be part of one representational system, or many.

A set of components all derived from the same source form a *description*. We use RS-descriptions to capture representational systems, which hold all components that any specific representation might be described using. A specific representation is captured through an R-description, which is primarily a subset of some RS-description. Problems are described with a Q-description, which is an R-description with an importance function, assigning each component a value between 0 and 1 based on how critical it is to capturing the problem.

Components and descriptions are intended to allow us to consistently catalogue problems, representations, and representational systems that are extremely

* Supported by the Hamilton Cambridge International Scholarship. Research conducted in conjunction with the rep2rep research team, who are supported by the EPSRC grants EP/R030650/1, EP/T019603/1, EP/R030642/1, and EP/T019034/1.

diverse. We worked to ensure that the framework followed three principles: formal and informal representations and systems were equally supported; no modality (sentential, diagrammatic, or otherwise) was favoured above any other; and the structure of the representation or system could be encoded as loosely or strictly as necessary. This ensures our framework is widely applicable, whether in our specific domain of interest—mathematics education—or more broadly.

Capturing similarities. The primary contribution of this PhD project is the theory of *correspondences*. A correspondence is a means of linking components from different RS-descriptions based on the probability of the component being present in R-descriptions of representations. This allows us to understand how information gets re-represented, and understand how effectively the informational content is preserved.

Correspondences are triples $\langle a, b, s \rangle$ consisting of a *source* component formula, a *target* component formula, and the correspondence's *strength* [5]. We consider the source component formula and the target component formula together to express a relationship: that these two things 'fill the same role' in their respective representational systems. For representations we say that a correspondence aims to capture how information can be re-encoded across representations. One caveat of this relationship is that it is not 'all-or-nothing': the information can be preserved to a certain degree, ranging from perfect down to completely lost. The degree of similarity is captured by the *strength* of the correspondence. Strength is defined in terms of the probability of an R-description containing a set of components that satisfy the component formula, and can be derived from a dataset of R-descriptions. The strength of the correspondence $\langle a, b, s \rangle$ is

$$s = \frac{\Pr(b \mid a) - \Pr(b)}{1 - \Pr(b)}.$$

Correspondences can be automatically derived in some circumstances [5].

To recommend a representational system, we define two measures through which their effectiveness can be algorithmically evaluated: informational suitability, and cognitive cost. Informational suitability considers the problem-specific aspects of a representational system: can it express everything the problem requires? Cognitive cost focuses on the human: is this representational system appropriate for this person to use? The complete definitions of these two measures are outside the scope of this summary, but described in [4].

We have produced an implementation of the properties and correspondence frameworks in a tool called robin. This implementation is fully automated from reading the descriptions through to making a recommendation based on the informational suitability calculation.

Evaluating the framework. To explore both the nature and efficacy of our framework we have performed two evaluations. First we performed an ablation study, in which we consider the output of robin with and without certain aspects of components and correspondences enabled. Second we ran a user study in which we asked mathematics teachers to perform the same representational system recommendation task that robin performs; we examine the responses of the teachers and consider how these responses map to those of robin.

The ablation study pulled apart the factors influencing the informational suitability score, and determined the contribution of each. We did this by surveying mathematics and computer science researchers to evaluate the informational suitability of representational systems for a specific problem, then correlated these with the scores produced by our implementation of our framework. We concluded that component *importance* and correspondence *strength* are both important contributors to the final score, and both are required to produce a result that correlates significantly with human experts.

The user study involved presenting mathematics teachers with descriptions of representational systems and student personas, and asking them to evaluate each system on its suitability to solve each of five problems in the cases of no particular student, a novice student persona, and an expert student persona. We established that our participating teachers do not have a consistent view of representational system suitability when controlling for problem and student profile, suggesting the task before us is more challenging than we first realised, and there may be factors for which we failed to control; what agreement could be extracted is consistent with the output of our implementation of the framework.

Contributions & future work. This work contributes a framework of *correspondences* that allow us to compute *informational suitability* and *cognitive cost* measures as a means of evaluating the effectiveness of representational systems. This work has wide applications, most immediately in intelligent tutoring systems: tools that can prompt students not just with the next step, but with hints tailored specifically to them which can reveal deeper insights and connections.

The next phase of research is to further generalise and formalise correspondences, with the aim to perform the transformation between representations automatically, similar to structure mapping [3]. We also hope to develop a suite of tools for the encoding and exploration of representational systems within our framework, meaning researchers and other experts can have a common set of tools to evaluate representational systems.

References.

1. Ainsworth, S.: The Educational Value of Multiple-representations when Learning Complex Scientific Concepts, chap. 9. Springer (2008).
2. Condell, J., Wade, J., Galway, L., McBride, M., Gormley, P., Brennan, J., Somasundram, T.: Problem solving techniques in cognitive science. Artificial Intelligence Review **34**(3), 221–234 (2010).
3. Gentner, D.: Structure-mapping: A theoretical framework for analogy. Cognitive Science **7**(2), 155–170 (1983).
4. Raggi, D., Stapleton, G., Stockdill, A., Jamnik, M., Garcia Garcia, G., Cheng, P.C.H.: How to (re)represent it? In: Alamaniotis, M., Pan, S. (eds.) IEEE 32nd International Conference on Tools with Artificial Intelligence, ICTAI. (2020).
5. Stockdill, A., Raggi, D., Jamnik, M., Garcia Garcia, G., Sutherland, H.E.A., Cheng, P.C.H., Sarkar, A.: Correspondence-based analogies for choosing problem representations. In: Homer, M., Hermans, F., Tanimoto, S., Anslow, C. (eds.) IEEE Symposium on Visual Languages and Human-Centric Computing, VL/HCC. (2020).