REAL 2 Knowledge Graph: Purposes, Context and Possibilities

Reviews, overviews, syntheses, meta-analyses all require the study of much material and make reference to yet more in order to explain their interpretation. The REAL 2 Review of Participatory Action Research (PAR) is no exception. Its subject matters range from well-defined areas of health research and practice through to how these are mediated by social interactions and by the methods and approaches used in the research process. It aims to be of interest to anyone working on any of these matters, including researchers, practitioners and those who as patients or citizens have participated in the process. That means that the way material referred to in the review is organised needs to accommodate not just the range of subjects covered but also the many ways and purposes with which potential users approach it.

With this context in mind, the immediate aims of this knowledge graph are to:

- Produce an easy to use interactive on-line resource for readers to further explore material and other links featured in the review
- Give more structure and depth to the bibliography of resources reviewed
- Make it easier for readers to follow up references, other material and people relevant to work on specific areas of interest within the review
- Allow authors of the review to share additional notes on the resources reviewed
- Give due prominence to work used in the review which has originated from less privileged settings and which is frequently algorthymmed out of results by major search engines

The graph has had to grapple with and seek solutions to a number of challenges which bedevil both traditional information management and the computerised interpretation of complex, transdisciplinary and intersectional information environments. Existing systems in both are, given their own internal logics, based on seeking to classify and define. Their approach to differences of understanding or interpretation is to seek to harmonise, reconcile or control. By contrast, realist analysis, PAR and other qualitative research requires the identification and recognition of difference as a precondition of any empirically sound process of analysis, negotiation or reconciliation. Our graph attempts to contribute to such processes by:

- Making explicit an understanding of the component parts of the study and the complex relationships between them
- Recognising the flexibility of language as it is used within and between different disciplines and communities of practice
- Seeking to find ways to both express and accommodate this lack of standardisation without creating confusion and misunderstanding

One feature of this graph is that this work, ideally, is done collectively or by a small team consulting constantly with other participants in the research process. It thus can be seen as a participatory tool in its own right. Members of a community of practice or of a new transdisciplinary research team can use it to jointly shape the field they intend to explore.

Its attention to language and to the key terms being considered by the research process in effect contributes to a new taxonomy which, while respecting the necessary precision of vocabularies used in bio-medical research, can be usefully applied to and used by the far wider range of settings in which health research and practice is shaped by and in turn shapes the societies in which it is embedded.

Such a process, in its fullest form, would extend far beyond the boundaries of the current review and would include a host of other specialist and local knowledges around issues such as food, housing etc. This graph, or the method on which it is based, offers a proto-type with a potential for a variety of institutions or groups to adapt and use, individually or in broader collaborations. The fact that it is web based, fully open source, potentially machine readable and based on established metadata standards all make it a natural tool for supporting an open ecology of knowledge.

Which brings us to AI. The history of ICT has been replete with the development (and sometimes fall) of techno-giants but also a mass of decentralised co-operation and exchange. Our concept of the knowledge graph is intended as a contemporary contribution to the latter. It does not engage with the language models which create AI, but offers a different, distributed and more carefully curated path to the collection and organisation of data – in our case metadata – on which the language models rely to produce good results. This, we suggest, is a strategic issue for anyone interested in research as a public good.

The graph as it stands is still in development and, no doubt, contains errors and requires further improvement. It's prime aim is to be of service to the REAL 2 programme as it seeks to publicise and explain its results. The aim of this concept note is to draw attention to the importance of the often hidden issues in the informational and computational contexts in which it has been developed.

Mike Powell
Emergent Works
June 2026