Why Visualization is an AI-Complete Problem (and why that matters)

Randy Goebel Alberta Innovates Centre for Machine Learning Department of Computing Science University of Alberta Edmonton, Alberta, Canada rgoebel@ualberta.ca

Abstract—Artificial Intelligence (AI) has infiltrated almost every scientific and social endeavour, including everything from medical research to the sociology of crowd control. But the foundation of AI continues to be based on digital representations of knowledge, and computational reasoning therewith. Because so much of modern knowledge infrastructure and social behaviour is connected to AI, understanding the role of AI in each such endeavour not only helps accelerate progress in those fields in which it applies, but also creates the challenges to extend the foundation for modern AI methods.

The simple hypothesis herein is that so-called AI-complete problems have a role in helping to articulate the appropriate integration of AI within other disciplines. With the current growth of interest in "big data" and visualization, we argue that relatively simple formal structures provide a basis for the claim that visualization is an AI-complete problem.

The value of confirming this claim is largely to encourage stronger formalizations of the visualization process in terms of the AI foundations of representation and reasoning. This connection will help ensure that relevant components of AI are appropriately applied and integrated, to provide value for a basis of a theory of visualization. The sketch of this claim here is based on the simple idea that visualization is an abstraction process, and that abstractions from partial information, however voluminous, directly confronts the non monotonic reasoning challenge; thus the need for caution in engineering visualization systems without carefully considering the consequences of visual abstraction. This is particularly important with interactive visualization, which has recently formed the basis for such fields as visual analytics.

Keywords-AI-complete visualization incomplete knowledge

I. INTRODUCTION

The process of visualization is largely about using computer programs to create visual abstractions of multidimensional data, with the general goal of guiding humans to a variety of inferences otherwise not obtainable from direct inspection of those data. We refer to *base data*, which is some collection of data to be rendered in a visual space, and *pictures*, including 2-, 3- and 4-dimensional pictures, as the result of transforming base data to a picture domain in which the human visual system can draw inferences.

Our informal argument about the AI-completeness of the visualization process is not one that becomes immersed in abstract formalization, but rather exposes connections across a relatively broad spectrum of formal philosophy, especially

formal language and representation, to make a connection between visual and logical expressions of information.

The chain of reasoning will not (yet?) withstand the precise scrutiny of a logician or complexity analyst, but the motivation is to make an initial connection, and help the growing community of visualization system engineers appreciate how pictures are really formal statements in an abstract visual formal language, and thus subject to all the same values and challenges of all formal representation and reasoning literature.

The argument proceeds by noting that visualization is about abstractions from base data to picture, that base data is largely incomplete, and that interaction with visualizations (so-called visual analytics) is akin to formal systems of representation used in AI question answering systems. The consequence is that interaction necessarily gives rise to non-monotonic reasoning. That consequence means that the identification of scope and content of context provides a direct connection to AI-complete problems.

The rest of this paper is organized as follows. The next section provides a brief summary of some of the origins of the idea of AI-completeness, and acknowledges past work on making the connection to traditional complexity measures of NP-complete and NP-hard. This connection is largely motivated by the desire to exploit traditional methods of Blum and others to characterize complexity classes, but adding a human in the loop. Following that, is an equally brief sketch of how visual analytics and visual interaction is analogous, perhaps even synonomous with the general idea of formal systems and interactive question-answering. That connection shows how even the simplest examples of historical logic puzzles and cognitive visual anomalies are closely related. This connection provides the basis for noting that the foundational concepts of compositionality and context determination arise within existing visualization system architectures, and must be at least part of the foundation of emerging theories of visualization.

Finally, should one believe the sketch of the proposition about the relationship between visualization and formal theories of context and compositionally, the consequent practical impact on visualization theories and systems is described, followed by a summary.

II. WHAT DOES IT MEAN TO BE "AI COMPLETE?"

The idea of AI completeness is attributed to Fanya Montalvo, but as Yampolskiy notes [1], the idea of AIcompleteness" has been a part of the field for many years and has been frequently brought up to express difficulty of a specific problem." Yampolskiy elaborates that position, and moves towards a more precise characterization of AI-completeness by considering a formalization of a human oracle embedded within a problem-solving context. The foundation concept is that of a problem being"Human Oracle solvable," which means that the combination of a standard human and a programmed machine can solve the problem. This is also the baseline for an earlier more precise characterization by Shahaf et al. [2], whose foundation problem is based on a "Human assisted Turing Machine." Both Yampolskiy and Shahaf et alexploit a framework where a conventional classification of problem complexity entails a declaration of the shared problem-solving responsibility of human and machine. For the complexity traditionalist, these background ideas provide the basis for considering the equivalent of polynomial reductions of AI problems to AI-Complete problems.

For this brief paper, the definitions from Yampolskiy [1, p. 7] provide a simple vocabulary within which one can make a statement that visualization is an AI-Complete problem:

- **Definition 1**: A problem C is AI-Complete if it has two properties:
 - 1. It is in the set of AI problems (Human Oracle solvable).
 - 2. Any AI problem can be converted into C by some polynomial time algorithm.
- **Definition 2**: AI-Hard: A problem H is AI-Hard if and only if there is an AI-Complete problem C that is polynomial time Turing-reducible to H.
- **Definition 3**: AI-Easy: The complexity class AI-easy is the set of problems that are solvable in polynomial time by a deterministic Turing machine with an oracle for some AI problem. In other words, a problem X is AI-easy if and only if there exists some AI problem Y such that X is polynomial-time Turing reducible to Y. This means that given an oracle for Y, there exists an algorithm that solves X in polynomial time.

While Shahaf et al. provide a notation based on the conventional complexity measure definitions of Blum ([3]), the reduction argument considered here will not require such notational machinery. Instead, we provide an informal reduction of interactive visualization to the general problem of non-monotonic reasoning problem, which is easily AI-complete. Nevertheless, our summary argument could be notationally, albeit naively, written as:

visualization \prec any AI – Complete problem C

So while these definitions provide at least an informal bridge

from statements like "natural language understanding is AIcomplete," they are convolved with concepts like the services of a human oracle; still the idea is like traditional complexity theory, where a problem can be transformed to another problem in a fashion similar to the familiar reductions of complexity theory.

Others, including both Yampolskiy and Shahaf et al. (e.g., [2]) have sketched a framework for characterizing tough problems as AI-Complete, and their lists of tough problems are, perhaps unsurprisingly, focused on both representation complexity (or sometimes referred to as expressive power of a representation), and computational complexity, which is relatively easily characterized as inference complexity, as in the original characterizations of NP-completeness (e.g., [4],[5]).

To provide some idea of the landscape of problems, here is a list sampled from both [1] and [2]:

Turing Test The baseline of all AI-complete problems.

Natural Language Question Answering The general problem of building a general question answer system.

- **Computer Vision** The general problem of human quality visual understanding.
- **ESP problem** The general problem of collaborative labelling [6].

With this briefest of AI-completeness background, the next step is to expose a difficulty technical problem that lies at the heart of interactive visualization. Once that connection is established, the informal reduction to existing AI-Complete problems will be, at least conceptually and informally, relatively straightforward.

III. INTERACTIVE VISUALIZATION AND THE PROBLEM OF NON-MONOTONICITY

The current research paradigm comprising "visual analytics" (e.g., [7], [8], [9] [10], [11], [12]) is very broad, but one can identify the following informal intersection of common properties:

- I Visualizations are abstractions of base data rendered in a visual domain Leaving aside debate about levels of abstraction, appropriate components of data visualization pipelines, how data is sampled, compressed, etc., there is general agreement that the mandate of any process of visualization is to transform base data to pictures of some kind.
- **II Visualizations provide guidance to preferred inferences** In general, the process of visualization is intended to expose inferences not otherwise easily made on base data. The choices made in the multi-level dimensionality reduction of data to 3D or 4D are generally made to encourage inferences on those data that are somehow more obvious to the human visual system.

III Interactive Visualization helps to amplify human visual inference Like adjusting numbers in

(a) $\exists x \ sloop(x) \land wants(x, John)$

(b) $\exists x \ sloop(x)$ $\exists x \ sloop(x)$ $\land wants(x, John)$ $\land name(x, "SloopJohnB")$ $\land wants(x, John)$

Figure 1. Resolution of the "John wants a sloop" ambiguity

a spreadsheet, human interaction with complex visualizations should provide the basis for elaborating obvious visual inference, for finding explanations underlying such inferences, and for speculative exploration of new hypothetical inferences.

A significant contribution to the complexity of visualization, even as abstracted in principles I, I, III above, is that, while almost all base data are somehow incomplete or noisy, visual transformations are still abstractions that somehow result from a process of dimensionality reduction. So even the simplest transformations from a set of base data to a visual representation can create at least ambiguity and often errors (e.g., [13]).

So what contributes to the claim that visualization AIcomplete? Part of what is required is to consider a picture as a formal collection of visual "sentences," like a set of logical statements, and to consider what happens in the context of an interaction with a human. Note here the role of the old idea of question-answering and interaction with a formal, at least digital, representation of information as some kind of representation from which inferences can be algorithmically drawn.

At least since the days of McCarthy's advice taker $[14]^1$, the simple idea behind man-machine interaction was that a machine's partial representation of the world could change as the result of interaction.

While there is a substantial and growing literature on interactive visualization and visual "analytics," (e.g., [15],[7]), the challenge to a formal semantical basis for visualization is substantial, if for no other reason that the repertoire of representations are literally all that humans can see (cf. formal language foundations of logical semantics).

So in the wide breadth of potential human interaction with visualizations, there needs to be some discussion about what a simple visual interaction might be, with the specific goal of showing how the interaction provides information to augment any initial visual representation.

To consider how conventional question answering and visual question-answering are connected, consider the following. Figures 1 and 2 may seem quite dissimilar, but will help illustrate how simple interaction in two different representation domains can both result in what McCarthy

¹originally published in 1959 and available at http://www-formal.stanford.edu/jmc/mcc59.pdf

anticipated as the result of a question-answering program accumulating new knowledge.

Figure 1 derives from Quine's old representation puzzle, where a simple logical representation of the phrase "John wants a sloop" can be rendered as

$$\exists X \ sloop(X) \land wants(X, John)$$

But as many formal philosophers have pointed out, this representation of the statement is ambiguous. As Quine so cleverly wrote, this rendering doesn't capture whether John wants a particular sloop, or mere relief from "slooplessness" [16, p. 177].

But without contesting the scholarly history of the formal philosophy of quantifiers and propositional attitudes, we can still see, as in Figure 1, a user interaction might resolve whether John has a particular sloop in mind, and resolve the ambiguity in expression Figure 1(a) as either the leftmost statement in 1(b), in which case the interaction would have indicated John is not fussy about which particular sloop he desires, or as in the rightmost statement in 1(b), which, in this case, declares John's interested in a particular sloop (in this case from the old folk song "The Sloop JohnB").



Figure 2. Resolution of the Necker cube ambiguity

Now returning to the visual ambiguity of the Necker cube, consider the diagram in Figure 2, taken from [17, p. 933]. As rendered in Figure 2(a), it is easy to consider a visual query which provides the interaction possibility to bias the visual interpretation to one of those in Figure 2(b), by providing the human user with the interaction capability of inserting a stick or similar cylindrical object into the "wire frame" Necker cube, thus showing which vertex is visually foremost (one of the two in Figure 2(b)).

This connection between logical and pictorial question answering will provide the basis for the argument that visualization is AI-Complete.

IV. COMPOSITIONALITY AND CONTEXT

The final component for our line of reasoning to confirm visualization as AI-complete continues to develop a parallel to existing paradigms of formal representation and reasoning, based on the assumption that pictures, even as abstractions of base data, can still be considered as sentences of a formal language.

In general, significant semantic - and ultimately computational or syntactic - challenges arise from the notion of compositionality. The general notion of compositionality assumes the preservation of component semantics as long as that semantics is context independent. In the field of formalisms for natural language processing, for example, this is *not* the case, and the formal approaches, typified by the legacy of Montague, has developed to ensure that context is a dynamic variable in semantical characterizations of compositional semantics. As well articulated by people like Hall-Partee [18] and Steedman [19], the Montaguelike pursuit of compositionally seeks to render ambiguity as unfulfilled context resolution, so that the ambiguity of the statement in Figure 1(a) might be directly captured by the introduction of an explicit context object, which can be composed within a logical expression something like

$wants(\gamma(context_{\alpha}, "John"), \gamma(context_{\alpha}, \exists x \ sloop(x)))$

This creates the opportunity for a compositional interpretation of context: whether John wants a particular sloop in context $context_{\alpha}$ is determined only with further information — we suggest user interaction to resolve the ambiguity as described above.

Here the important point is that determination of context, whether in resolving ambiguity in logical statements or pictures, is potentially NP-complete from the logical inference point of view (simply consider enumerating all possible models for any ambiguous sentence). To find the analog of context in a visual space, consider again the diagram of Figure 2, where one of two possible interpretations is provided by inserting a stick in the 3D interpretation of the Necker cube. In more complex scenes, the potential number of "sticks" may well be as large as the number of models, for a suitably expressive logical formalism.

The last component necessary for the reduction argument notes the potential non-monotonicity of interactive visualization, for example, as noted in Goebel et al. [20]. Goebel et al. demonstrate how a naive interpretation of a visual "pinch" operation on a classical visualization (see Figure reffig:minard) can create ambiguous consequences because the base data is simply incomplete. Consider that a "pinch" at the designated point in Figure reffig:minard has no single interpretation based on the base data, because the correlation with pinching and reduction in the size of Napoleon's army is simply not given. The direct consequence is that any visualization interaction must be directly constrained by the



Figure 3. Ambiguity of a "pinch" operation

semantics of the underlying data, *not* the visual syntax of the interaction (e.g., pinching your photograph on a tablet doesn't make you smaller).

Now the conventional reducibility, approached within a framework of an AI-completeness formalization, can be connected in several ways: this includes a formalization of the complexity of determining context, which is at least exponential in the size of context variables. In the case of logical representations, that complexity is related to the expressiveness of the underlying formal language (e.g., propositional logic, first order logic with unary predicates, etc.). In the case of visualization, the characterization is much more difficult: it is related to the how a given repertoire of interactive manipulations can affect the values of that range of visual variables which can be adjusted by any such interaction (cf. Figure 3).

V. WHY IT MATTERS: IMMEDIATE CONSEQUENCES FOR INTERACTIVE VISUALIZATION

The argument that interactive visualization is AI-Complete can be repeatedly refined with more detail and precision, like that of the framework of Shahaf et al. [2] summarized above. But that pursuit would only address and perhaps reduce the skepticism on this informal argument.

Another avenue of discussion that may better help connect conventional visualization research is to consider a visualization framework where the base data is unavailable. For example, the ReVision system of Savva et al. ([21]) takes a novel approach to improving visualizations by first creating a fairly accurate classification and analysis of a set of standard visualization formats (e.g., pie chart, histogram, radar plots, etc.). Assuming no access to the base data that gave rise to a visualization, the ReVision system then employs an inductive/machine learning approach to accurately reconstruct that data conveyed in any one such instance. Once that process completes, the reconstructed data can then be re-rendered into an improved visualization.

So the inductive process that, for example, extracts the scale marks from a pie chart, creates the basis for transforming a less than effective rendering into a perhaps more perspicuous visualization, for example, from a pie chart to a histogram, as in Figure 4.

Note that ReVision, having no access to the base data from which the visualization was abstracted, must induce that data required to recreate an alternative visual format. The idea that this inductive process can wrongly induce any portion of the intended base data of a visualization is exactly where the incompleteness of accessible base data coincide. The consequence is that the accuracy of any ReVision transformation from one chart style to another depends on the accuracy of the reconstruction of base data.

In the more general case, as illustrated by the Minard example, the arbitrary completion of missing base data produces a bias in any choice of alternatives arising from visualization interaction.



Figure 4. ReVision transformation from pie chart to histogram

The practical consequences for visualization, and the engineering of visualization systems, are more important and they are generally twofold:

- 1) Any repertoire of visualization interactions must be coupled with semantic descriptions of the base domain.
- The consequences of visualization interactions on partial data, no matter how dimensionality reduction and abstraction is done, will be necessarily non-monotonic.

In general, the visualization interaction problem is at least as complex as the general non-monotonic reasoning problem [22]. This means, for example, that any visualization theory or architecture (e.g., [12], [15]) must include a component that addresses the general complexity of visual reasoning.

In addition, the considerable accumulation of tools and techniques for visual analytics (e.g., [23], [24]) will have to consider how to couple visualization constraints in any general visualization architecture, in order to avoid creating the equivalent of visual anomalies from any abstraction of base data in picture space.

This will require the deeper understanding and the disciplined engineering of emerging formal characterizations of visual anomalies, like those considered by Sughihara [25]) and Mortensen ([26]).

VI. SUMMARY AND CONCLUSIONS

If there is value in considering the classification of problems in the category of AI-Complete problems, one must consider at least some detail in the formulation of AIcompleteness (as have [2], [1]), and then focus on what any reductions of existing problems suggest about the challenge for AI tools and techniques.

Here, we suggest that the process of interactive visualization is both an abstraction process and a reasoning process, and thus has a close correspondence with formal representations of knowledge and formal methods of reasoning. With this correspondence made, the AI-completeness of interactive visualization (or, for some, "visual analytics") emerges from the general problem of reasoning with incomplete information, and the general challenge of non-monotonic reasoning.

The immediate consequences are positive and negative: the negative are that all visualization architectures must consider the complexity of reasoning from incomplete information, and must be careful to ensure that any repertoire of visual interactions are semantically coupled with the data domains in question. The positive consequences are that there are emerging theories of how to formally and precisely avoid the creation of visual anomalies, which can be incorporated into theories of interactive visualization.

ACKNOWLEDGEMENTS

This research is supported by the Canadian Natural Sciences and Engineering Research Council (NSERC) and the Alberta Innovates Centre for Machine Learning (AICML).

REFERENCES

- R. Yampolskiy, "Turing test as a defining feature of AI-Completeness," in Artificial Intelligence, Evolutionary Computing and Metaheuristics (X.-S. Yang, ed.), vol. 427 of Studies in Computational Intelligence, pp. 3–17, Springer Berlin Heidelberg, 2013.
- [2] D. Shahaf and E. Amir, "Towards a theory of AI completeness.," in AAAI Spring Symposium: Logical Formalizations of Commonsense Reasoning, pp. 150–155, AAAI, 2007.
- [3] M. Blum, "A machine independent theory of the complexity of recursive functions," *Journal of the ACM*, vol. 14, no. 2, pp. 322–336, 1967.
- [4] J. Edmonds, "Paths, trees, and flowers," Canadian Journal of Mathematics, vol. 17, pp. 449–467, 1965.
- [5] S. A. Cook, "The complexity of theorem-proving procedures," in *Proceedings of the Third Annual ACM Symposium on Theory of Computing*, STOC '71, (New York, NY, USA), pp. 151–158, ACM, 1971.
- [6] L. von Ahn, "Games with a purpose," Computer, vol. 39, no. 6, pp. 92–94, 2006.
- [7] D. A. Keim, F. Mansmann, J. Schneidewind, J. Thomas, and H. Ziegler, "Visual analytics: Scope and challenges," in Simoff *et al.* [15], pp. 76–90.
- [8] D. Keim, G. Andrienko, J.-D. Fekete, C. Görg, J. Kohlhammer, and G. Melançon, "Visual analytics: Definition, process, and challenges," in *Information Visualization, Lecture Notes* in Computer Science 4950, pp. 154–175, Springer, 2008.

- [9] S. Bertschi, S. Bresciani, T. Crawford, R. Goebel, W. Kienreich, M. Lindner, V. Sabol, and A. V. Moere, "What is knowledge visualization? perspectives on an emerging discipline," 2010 14th International Conference Information Visualisation, vol. 0, pp. 329–336, 2011.
- [10] M. Chuah and S. Roth, "On the semantics of interactive visualizations," in *Information Visualization '96, Proceedings IEEE Symposium on*, pp. 29–36, Oct 1996.
- [11] R. Goebel, W. Shi, and Y. Tanaka, "The role of direct manipulation of visualizations in the development and use of multi-level knowledge models," in *Proceedings of the* 17th International Conference on Information Visualisation, pp. 325–332, IEEE, 2013.
- [12] R. Patterson, K. Liggett, L. Blaha, G. Grinstein, P. Havig, D. Kaveney, J. Moore, and K. Sheldon, "A theory and model of visualization," *International Journal of Computer Graphics* and Visualization, 2013. (preprint).
- [13] S. Kaski and J. Peltonen, "Dimensionality reduction for data visualization," *IEEE Signal Processing Magazine*, pp. 100– 104, 2011.
- [14] J. McCarthy, "Programs with common sense," in *Semantic Information Processing*, pp. 403–418, MIT Press, 1968.
- [15] S. J. Simoff, M. H. Böhlen, and A. Mazeika, eds., Visual Data Mining - Theory, Techniques and Tools for Visual Analytics, vol. 4404 of Lecture Notes in Computer Science. Springer, 2008.
- [16] W. Quine, "Quantifiers and propositional attitudes," *The Journal of Philosophy*, vol. 53, no. 5, pp. 17–187, 1956.
- [17] P. Klink, R. van Wezel, and R. van Ee, "United we sense, divided we fail: context-driven perception of ambiguous visual stimuli," *Philosophical Transactions of the Royal Society B*, vol. 367, no. 7, pp. 932–941, 2012.
- [18] B. Hall-Partee, "On the history of the question of whether natural language is"illogical"," *History* and *Philosophy of the Language Sciences*, 2013. http://hiphilangsci.net/2013/05/01/on-the-history-of-thequestion-of-whether-natural-language-is-illogical.
- [19] M. Steedman, "Semantics for semantic parsing," ACL Workshop on Semantic Parsing, 2014. http://yoavartzi.com/sp14/slides/steedman.sp14.pdf.
- [20] R. Goebel, "The challenge of semantic symmetry in visualization," in *Proceedings of the 18th International Conference* on Information Visualisation, pp. 272–33, IEEE, 2014.
- [21] M. Savva, N. Kong, A. Chhajta, L. Fei-Fei, M. Agrawala, and J. Heer, "Revision: Automated classification, analysis and redesign of chart images," in *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*, UIST '11, (New York, NY, USA), pp. 393–402, ACM, 2011.
- [22] M. Cadoli and M. Schaerf, "A survey of complexity results for non-monotonic logics," *Journal of Logic Programming*, vol. 17, pp. 127–160, 1993.

- [23] R. Spence, *Information Visualization: Design for Interaction*. Prentice-Hall, 2007.
- [24] R. Spence, Information Visualization: An Introduction, Third Edition. Springer, 2014.
- [25] K. Sugihara, "Design of solids for antigravity motion illusion," *Computational Geometry*, vol. 47, pp. 675–682, 2014.
- [26] C. Mortensen, *Inconsistent Geometry*. Studies in Logic, College Publications, 2014.