Toward Diagrammatic Automated Discovery in Axiomatic Physics*

Atriya Sen[1,4] • Selmer Bringsjord[1,2,3,4] • Nick Marton[1,4] • John Licato[1,4]
Department of Computer Science[1]
Department of Cognitive Science[2]
Lally School of Management & Technology[3]
Rensselaer AI & Reasoning (RAIR) Laboratory[4]
Rensselaer Polytechnic Institute (RPI)
Troy NY 12180 USA
Contact: Sen (`sena2@rpi.edu`)

preliminary draft of 040415

Abstract

At the First International Conference on Logic and Relativity, Bringsjord, on behalf of a trio of RAIR-Lab researchers, showed a formal, semi-automated, symbolic proof of Theorem NEAT (No Event at Two Places). Extension and refinement of this research appeared subsequently (in *Synthese*). This prior work, like 99.9% of proof-oriented work in the formal sciences, is *homogeneously linguistic* in nature: the proofs in question are based exclusively on formal *languages*; diagrams, pictures, images, etc. are nowhere to be seen. Yet mathematical physicists routinely employ (informal) visual and diagrammatic reasoning in their proofs. A formal system leveraging *both* visual and symbolic reasoning enables *heterogeneous* proofs that are (i) not only more readable, intuitive, and consistent with scientific practice, but also (ii) simpler (in a formal sense), and therefore potentially easier for machines to discover on their own. Herein, we announce the availability of precisely such a system, one built directly atop Vivid, a heterogeneous logicist framework in turn built atop **denotational proof languages** (DPLs); and we employ the system to move closer to a formal, semi-automated proof of Theorem NEAT that is at once both linguistic and diagrammatic.

*We are indebted to Naveen Sundar G. of Yahoo for substantial work in computational logicist physics, carried out while in Bringsjord's RAIR Lab. This paper owes a large debt to Konstantine Arkoudas, whose seminal DPL/Vivid research plays a pivotal role. Finally, without the generous support of AFOSR, Bringsjord's move into computational logicist physics would have been flat-out impossible.

Contents

List of Figures

1 Introduction; Plan

"We do not listen with the best regard to the verses of a man who is only a poet, nor to his problems if he is only an algebraist; but if a man is at once acquainted with the geometric foundation of things and with their festal splendor, his poetry is exact and his arithmetic musical."
- Ralph Waldo Emerson, *Society and Solitude* (1876)

At the First International Conference on Logic and Relativity, one of the authors (SB) presented "Proof Verification and Proof Discovery for Relativity," wherein was described the semi-automated discovery (and verification) of a proof for Theorem NEAT (No Event At Two Places), using the proof-engineering environment Slate (Bringsjord, Taylor, Shilliday, Clark & Arkoudas 2008), a fertile system for RAIR-Lab formal, proof-based AI (e.g. see Bringsjord, Govindarajulu, Ellis, McCarty & Licato 2014).[1] This work was subsequently refined and expanded, and published in *Synthese* (Govindarajalulu, Bringsjord & Taylor 2014). In that paper, a future-work agenda was adumbrated, and the authors specifically mentioned that (i) informal physics proofs often draw upon visual, geometric concepts not directly present in standard formal languages (such as the formal languages that anchor first-order logic, higher-order logic, modal logic, etc.), such as lines, planes, slopes, et cetera, and that (ii) the reasoning propounded is commonly heterogeneous in nature: it mixes diagrammatic reasoning on the one hand, with informal linguistic/sentential reasoning naturally modeled in standard formal languages on the other hand.

In the time since that conference, a first-order proof of the theorem "no inertial observer can travel faster than light" has also been verified in the proof assistant and interactive theorem prover Isabelle/HOL by Stannett and Németi (2014). Their proof models the geometric structures common to all models of spacetime, such as vectors, points, lines and planes. However, it is important to note that these models are *linguistic* (or sentential) in nature, and therefore the reasoning involved remains classically homogeneous: none of the diagrams and pictures common in the presentation of such reasoning by one human to another are present in the formal reasoning in question.

The Vivid system (Arkoudas & Bringsjord 2009*b*) for mechanized heterogeneous natural deduction, in contrast, was developed explicitly to combine linguistic and diagrammatic representations, and deduction over such representations, at the formal level. The framework represents a member of the family of **denotational proof languages** (DPLs) invented by Arkoudas (2000). Vivid is able to: model under-determined diagrams, and deal with incomplete information; includes a formal semantic framework based on a Kleenean three-valued logic; includes general inference mechanisms for the valid extraction of information from diagrams and the incorporation of sentential information into diagrams; and has been proved sound.

Given this context, the sequence of the sequel is as follows: Section 2 is a brief, non-technical overview of Vivid, and of an implementation due to one of the authors (NM). Section 3 provides a diagrammatic proof of a special-relativity theorem that can be directly respresented in a Vivid language and then automatically verified. Section 4 describes PAGI World (Licato et al. Submitted for Initial Review), written in the Unity game engine, which we use to simulate a certain proper subset of the space of Vivid proofs. Section 5 wraps up the paper by among other things pointing toward our goal of Vivid-based proof-*generation*, to enable automated scientific discovery in axiomatic physics.

2 Vivid Encapsulated

Diagrams are arguably the most useful tool for visual inference (including deductive inference, our focus). As opposed to purely linguistic/sentential representations of information, diagrammatic representations enable visual reasoning. Such representations have structural correspondences with what they represent; these **analogical** representations (to use the terminology of Sloman 1971) genuinely resemble what they depict.[2] Often resolution of an image or incomplete information about the representation can make extracting sentential information from a scene impossible; this is not the case in diagrammatic representations. In the Vivid system, diagrammatic representations are held as partially complete descriptions

[1] Slate is based on natural deduction carried out in hypergraphs, a novel form of natural deduction.

[2] Other thinkers use different adjectives to denote the same kind of representations. E.g., Barwise & Etchemendy (1995) talk of **homomorphic** representations.

of finite system states; this gives the capability of reasoning with incomplete information (which is unavoidable) to any system that employs diagrammatic reasoning. This robust ability of diagrammatic representations to incorporate incomplete information into proofs provides a convenient and efficient solution to the challenge of formalizing and implementing reasoning that is part linguistic, and part diagrammatic. Accordingly, Vivid's proof theory includes sentential deductions (which D ranges over), diagrammatic deductions (which Δ ranges over), and heterogeneous deductions that combine these two types of deductions (which $\mathcal{D}$ ranges over). All deduction is **natural** deduction, rather than for instance resolution.[3]

In the following example, the use of Vivid gives rise to an extended concept of "four-dimensional proof."[4] First, we form a purely sentential representation to define the mathematical relations we want to incorporate in the visual scene. Then, we utilize PAGI World's sensors to receive purely visual information from the scene.[5] We then transform this visual information into a more convenient sentential representation, and then use this representation to form a diagrammatic representation of the scene itself, in which we solve the task at hand. Our example specifically uses Vivid's **thinning rule** (Arkoudas & Bringsjord 2009*a*) (see note 3). In this way, we go from a purely sentential representation, to a purely visual one, then from visual to sentential, and finally from sentential to visual.[6]

Consider the scene in Figure 11. In this scene, the agent is presented with two clocks. We obtain a particular instance of a Vivid proof by employing the suite of inference schemata in (Arkoudas & Bringsjord 2009*a*) to determine whether or not the first clock is ahead of the second. While the agent receives visual information from the scene, we restrict the sensors so that the agent can *roughly* determine the time on each clock: the agent receives a set of approximately correct times; his vision is a bit fuzzy. Then, to interpret the `Ahead` relation, we employ the thinning rule, trying each value in the sets of potential times of the clocks against each other. Since the first clock is definitively further ahead of the second clock in the agent's vision, the agent can deduce that the `Ahead` relation holds, and that the first clock is indeed ahead of the second.

3 Example

This section concerns Theorem 2.2 from "Logic of Space-time and Relativity Theory" (Andréka, Madarász & Németi 2007), which states that no observer observes the same event at two different space-time locations in models of the field and light axioms of special relativity. We abbreviate this, as before, as 'Theorem NEAT' (No Event At Two Places). An informal proof is supplied in English; the proof uses geometric constructs in its reasoning. We now describe how this proof may be formulated in a formal heterogeneous (i.e., linguistic/sentential-and-diagrammatic) manner, and represented in a Vivid language for automatic verification.

> **Argument**: Consider a 2-dimensional plane with two axes, the horizontal axis representing the spatial dimension and the vertical axis representing the temporal dimension. Let there be an inertial

[3]Due to space constraints, we cannot inform the reader about the many deductions available in the space $\mathcal{D}$. Some of the deductions of type D will be fundamentally quite familiar to all readers (e.g.,

$$\textbf{specialize } \forall x_1, x_2, \cdots, x_n.F \text{ with } t_1, \cdots, t_n)$$

while others are fully novel. E.g., "sentential-deduction" cognoscenti will not have encountered the present paper having already in hand an understanding of:

$$(\sigma; \rho) \text{ by } \textbf{thinning} \text{ with } F_1, \cdots, F_n.$$

We do provide some selected information below, to facilitate exposition.

[4]Such proofs are impossible in standard, linguistic logics. For example, **spatial logics**, while ironically targeting phenenomena that are often depicted with diagrams by humans, are exclusively linguistic. This is revealed e.g. by study of (Aiello, Pratt-Hartmann & Benthem 2007*b*). In fact, the purely linguistic nature of the logics in question is explicitly affirmed; e.g., we read:

> By a *spatial logic*, we understand any formal **language** interpreted over a class of structures featuring geometrical entities and relations, broadly construed. (Aiello et al. 2007*a*, p. 1; bolded text our emphasis)

[5]PAGI World, described below, provides both a realistic, physical environment and sufficiently rich visual information, and is thus a natural choice for an environment that supports heterogenous deduction specified in Vivid.

[6]Technically, deduction by **cases** is what constitutes the key nexus between linguistic/sentential and diagrammatic deduction in Vivid, but space in the current, preliminary version of the present paper prevents us from giving a full presentation/explanation. However, we do deploy **cases** below (§3) and doing so provides what we take to be an appreciable degree of explication of this kind of deduction.

observer m at the origin. Consider two distinct points x and y on the plane.

From the field axioms, it follows that through any point x, there will be two lines of slope 1, where the speed of light $c = 1$.

We now summarize the process of formulating a diagrammatic representation in Vivid, and provide such a formulation. To specify a formula instance in Vivid, three steps are necessary.

First, we must provide an Attribute Structure A containing a collection of attributes and the set of their possible values, along with a set of computable relations $\mathcal{R}$ involving those attributes. We choose the attributes `position`, `slope`, and `line_positions` with ranges $\mathbb{R}^n$, $[0,1]$, and $\mathbb{R}^{n^*}$ (where $*$ is taken to mean the set of all lists of the elements of $\mathbb{R}^n$), respectively. Furthermore, we choose the relations $R_1 \subseteq \mathbb{R}^n \times \mathbb{R}^{n^*}$ defined as $R_1(p_1, [p'_1, ..., p'_k]) \Leftrightarrow p_1 \in [p'_1, ..., p'_k]$, $R_2 \subseteq \mathbb{R}^n \times \mathbb{R}^n$ defined as $R_2(p_1, p_2) \Leftrightarrow p_1 = p_2$, and $R_3 \subseteq [0,1]$ defined as $R_3(s) \Leftrightarrow s = 1$. Then, in Vivid's native notation, our attribute structure is $A = (\{position : \mathbb{R}^n,\ slope : [0,1],\ line\_positions : \mathbb{R}^{n^*}\}; R_1, R_2, R_3)$, where R_1, R_2, and R_3 are defined as above.

Next, we must specify a Vocabulary $\Sigma = \mathbf{(C,R,V)}$, consisting of a set of constant symbols $\mathbf{C}$, a set of relation symbols $\mathbf{R}$, and a set of variables $\mathbf{V}$ used as the signature for the state we are representing. We define $\mathbf{C}$ as $\{x, y, m, l_1, l_2\}$; $\mathbf{R}$ as the set containing $through(p, l)$, which holds when point p lies on line l; $s(p, x)$, which holds when photon p is observed at point x; $slope\_of\_one(l)$, which holds when the slope of line l is 1; and $\mathbf{V}$ as $\emptyset$, i.e., the empty set.

Finally, we must provide an interpretation of the relation symbols of Σ into A. That is, as dictated by the formal machinery of Arkoudas & Bringsjord (2009*b*), we must provide a mapping I that assigns to each relation symbol $R \in \mathbf{R}$ of arity n:

1. a relation $R^I \in \mathcal{R}$ of some arity m, called the **realization** of R:

$$R^I \subset A_{i_1} \times \cdots \times \subset A_{i_k}$$

   (where we might have $m \neq n$); and

2. a list of m pairs

$$[(l_{i_1}, j_1), ..., (l_{i_m}, j_m)],$$

   the **profile** of R, denoted by $Prof(R)$, with $1 \leq j_x \leq n$ for each $x = 1, ..., m$.

We summarize this interpretation in the following table.

Symbol	**Arity**	**Realization**	**Profile**
through	2	R_1	$[(position, 1), (line\_positions, 2)]$
observes	2	R_2	$[(position, 1), (position, 2)]$
slope_of_one	2	R_3	$[(slope, 1)]$

We now proceed by cases. Vivid has a control-construct built in for the formalization of proofs by cases.

> **Case I** (1): y is not on either of these two lines. In this case, assume that m observes a photon p at x, and observes the same photon p at y. However, the line connecting x and y would have a slope other than 1, which implies that the photon would be traveling at a speed other than the speed of light. This contradicts the light axiom. It follows that it is impossible that m should observe the same event at x and y.
>
> **Case II** (2): y is on one of these lines. Consider the other line $l1$, also of slope 1, through x. Let there be a photon q somewhere on this line other than at x. This photon may legitimately have traveled from x to its current location in space-time; hence q represents an event m observes at x. Now assume that m observes the same event at y. This would mean that q is somewhere on line $l2$, the line of slope 1 through y which is other than the line passing through x. Then, q must be on both lines $l1$ and $l2$, and hence at the point where they meet. However, there is no such point, since $l1$ and $l2$ are parallel to each other. Hence, m cannot observe the event represented by q at y. It follows that it is impossible that m should observe the same event at x and y. **QED**

We now define an additional axiom from our relations defined in Σ above, and proceed to demonstrate our proof within the framework of Vivid. We define the *not_on_same_worldline axiom* as the Horn clause: $observes(p, x) \wedge through(x, l) \wedge \neg through(y, l) \rightarrow \neg observes(p, y)$, taken to mean informally

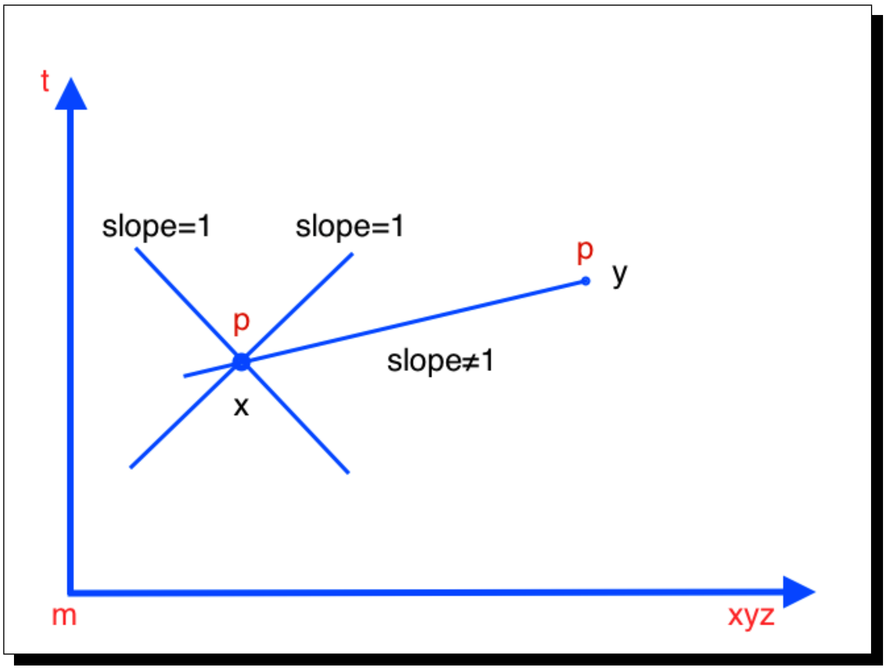

Figure 1: Case I

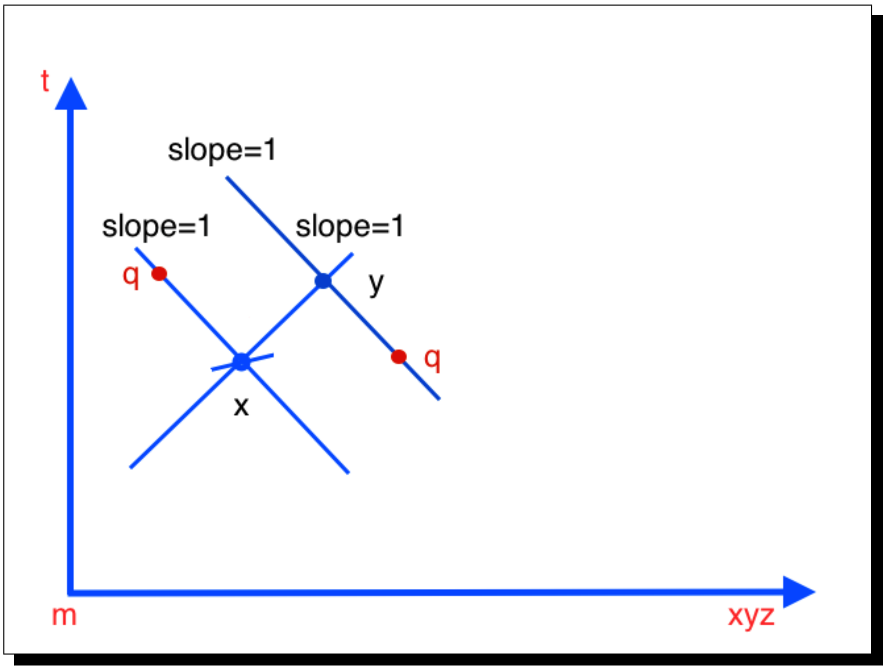

Figure 2: Case II

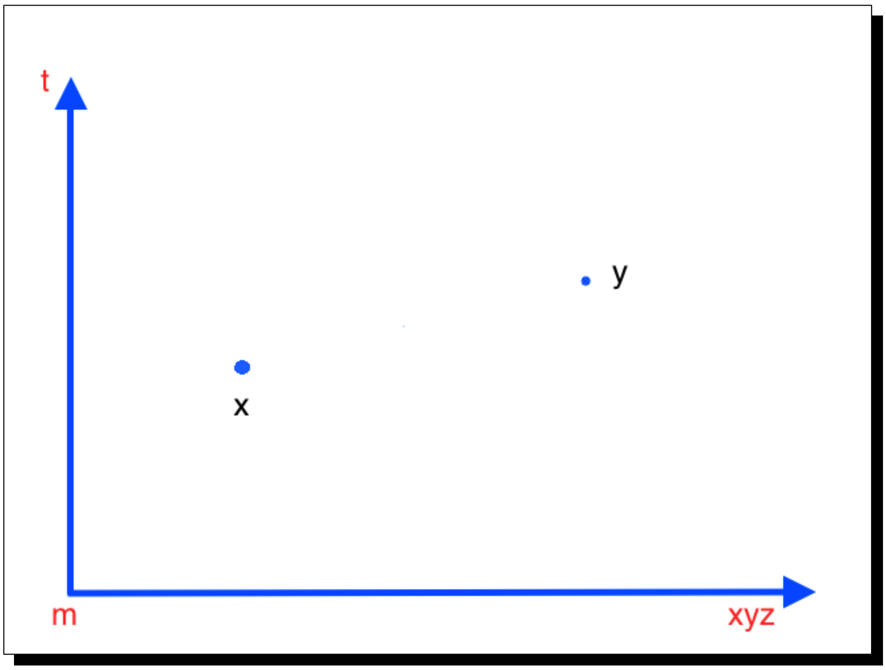

Figure 3: Δ_0

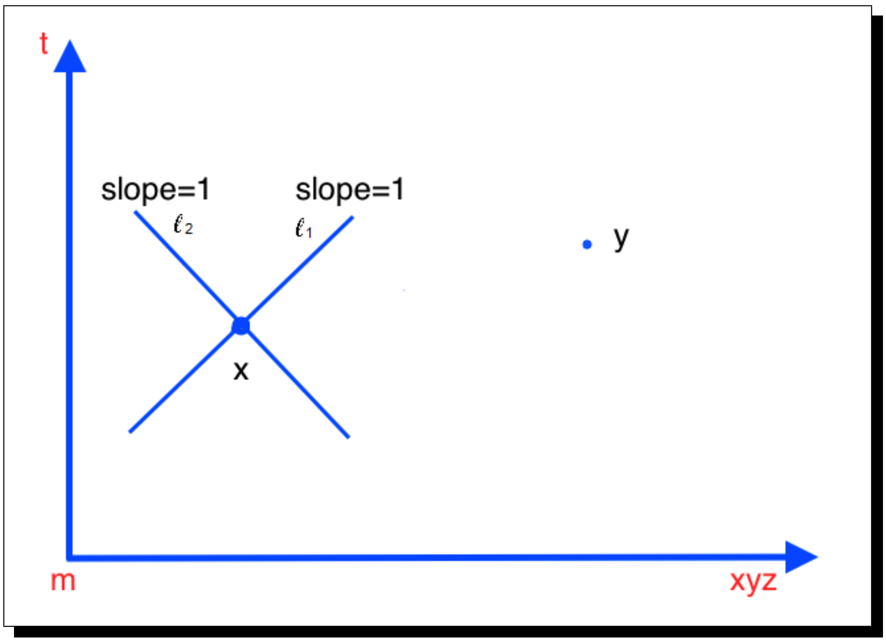

Figure 4: Δ_1

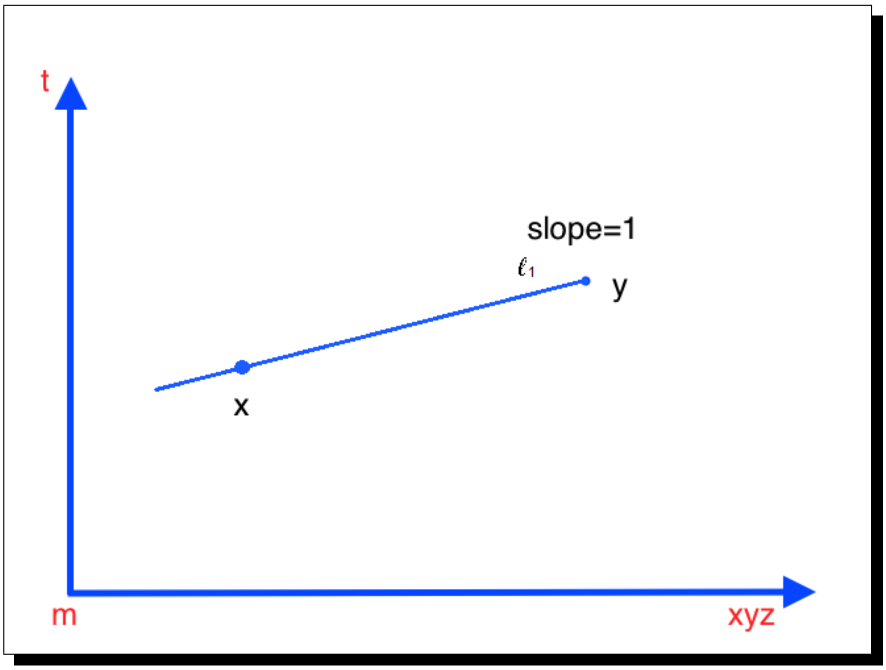

Figure 5: Δ_2

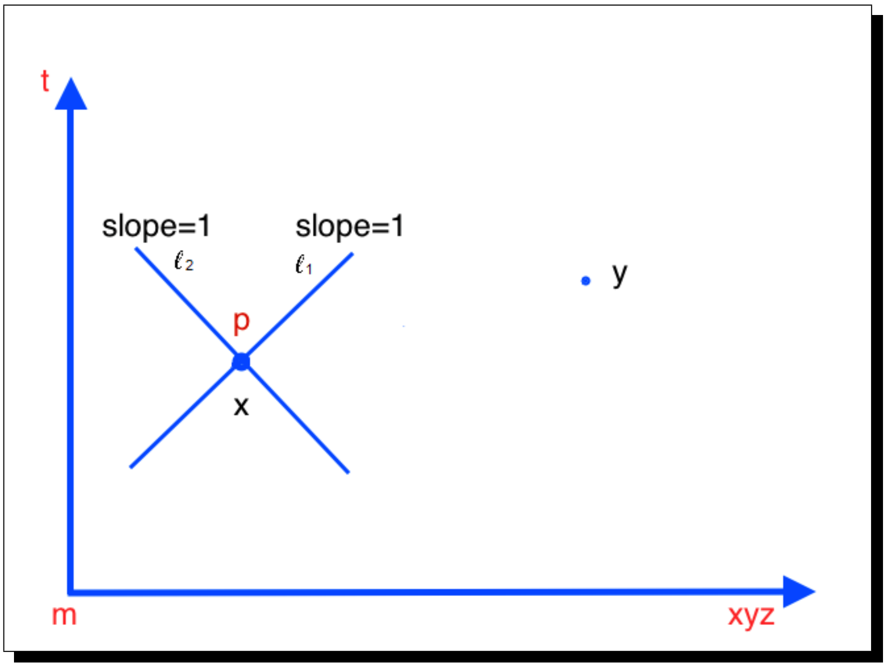

Figure 6: Δ_3

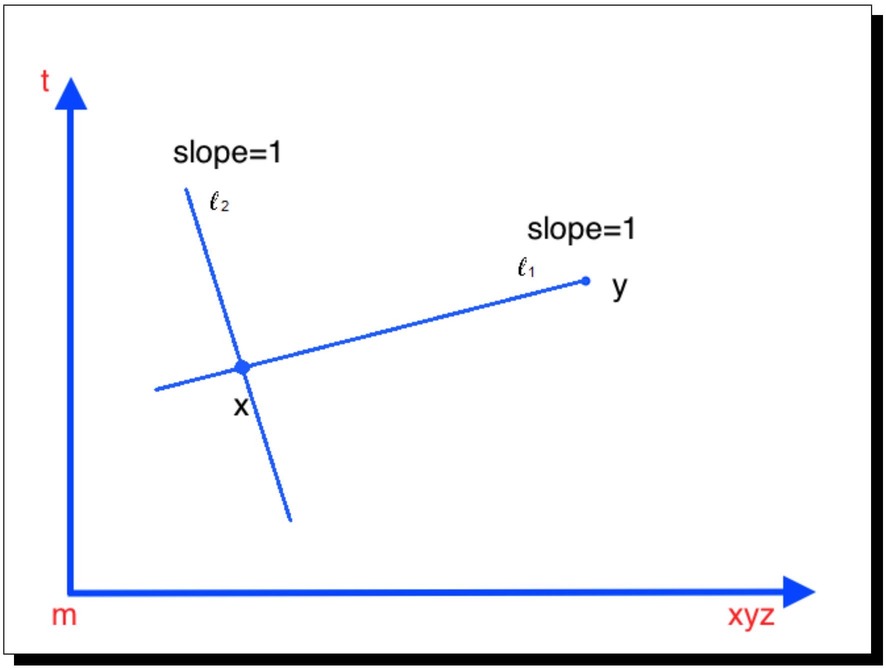

Figure 7: Δ_4

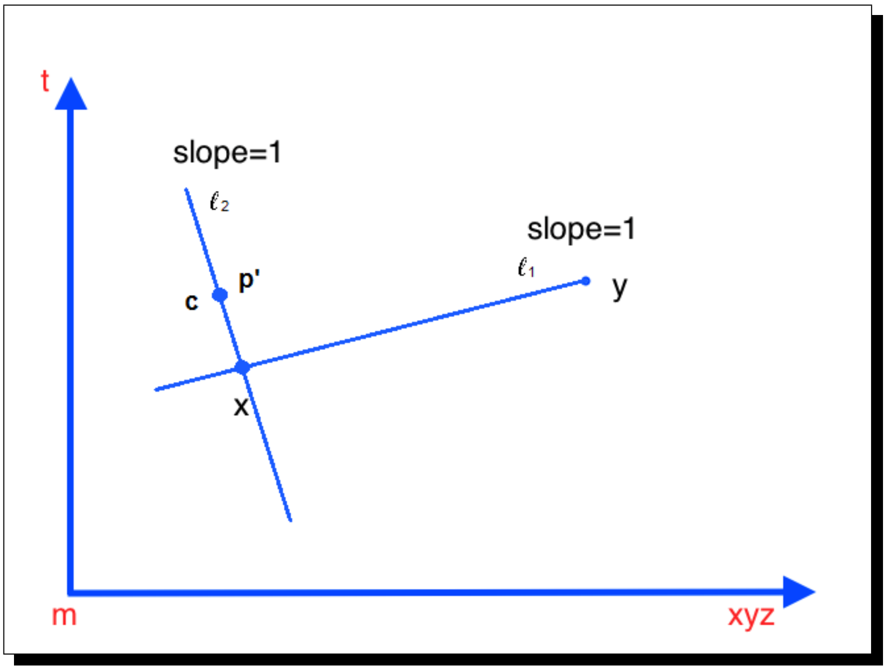

Figure 8: Δ_5

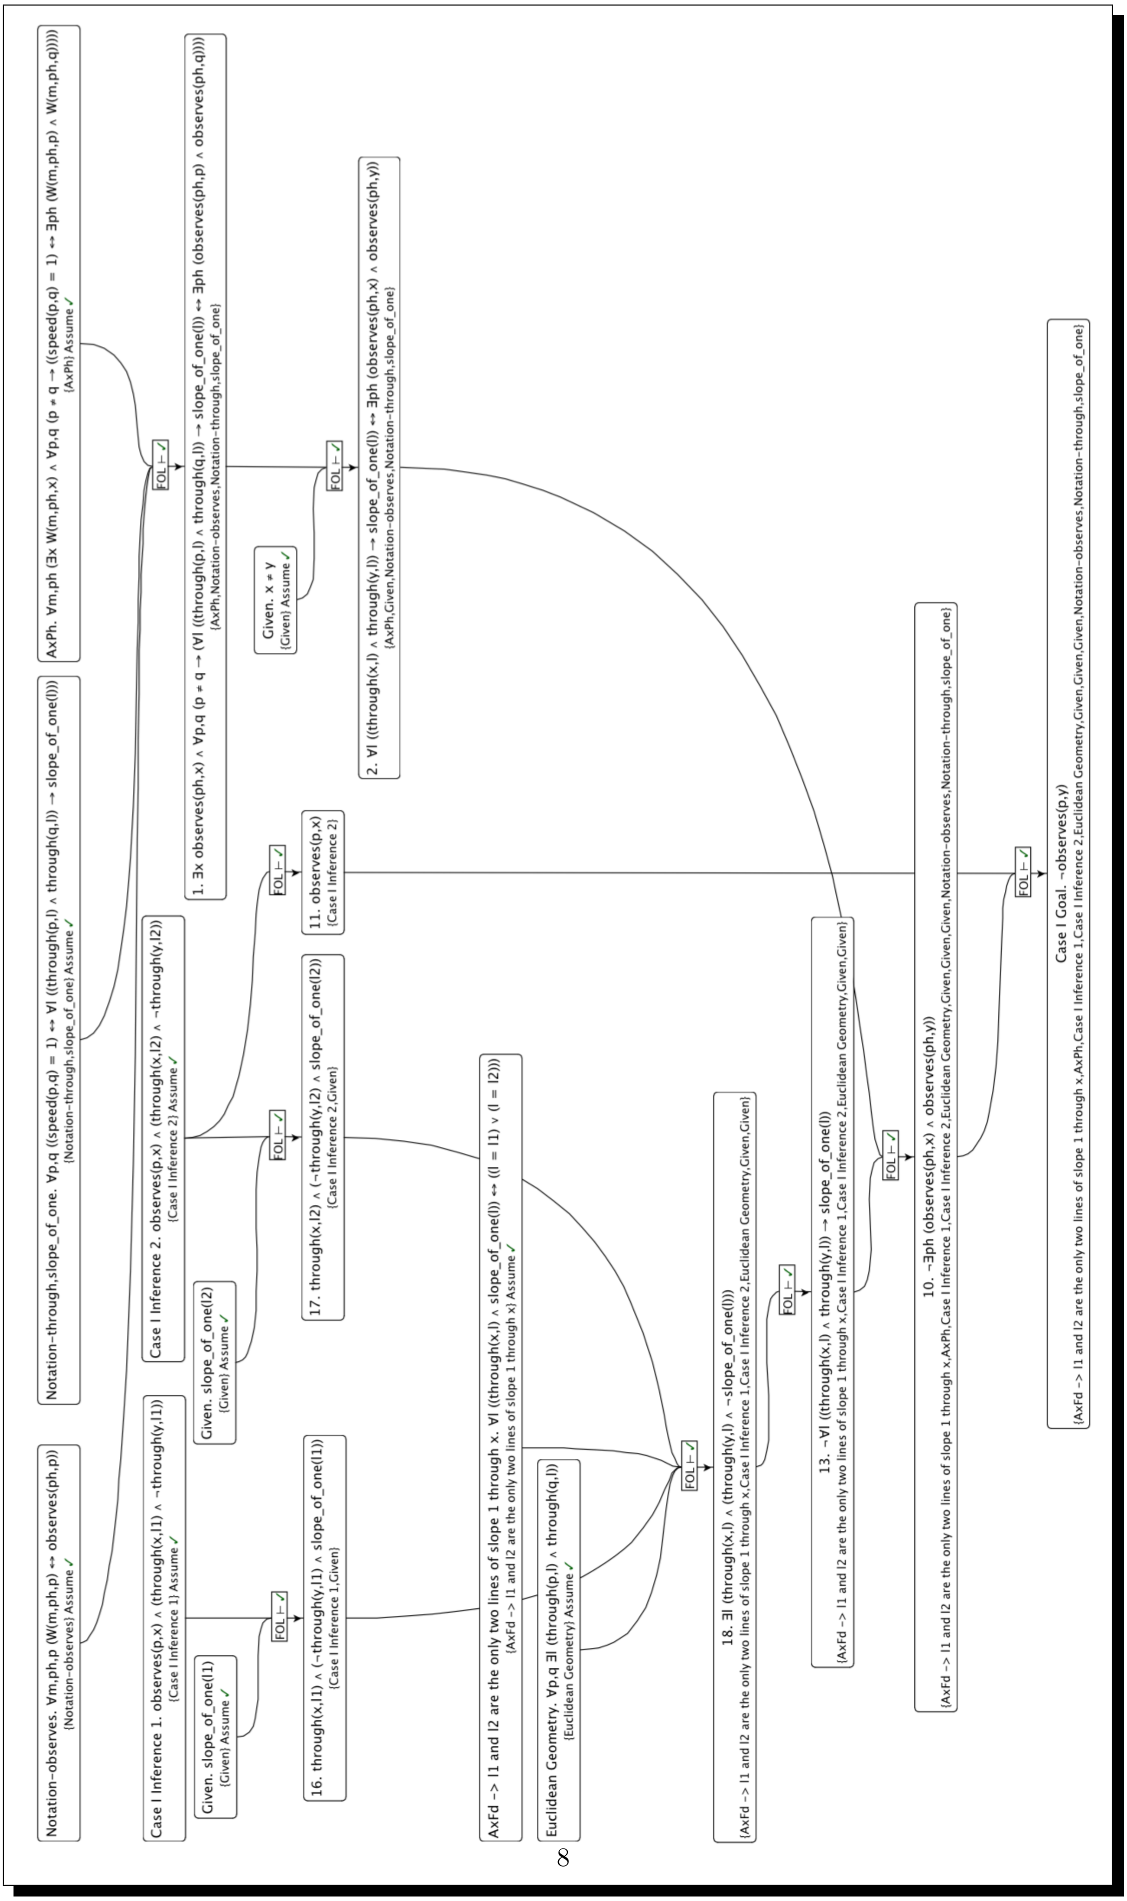

Figure 9: Proof for Case I

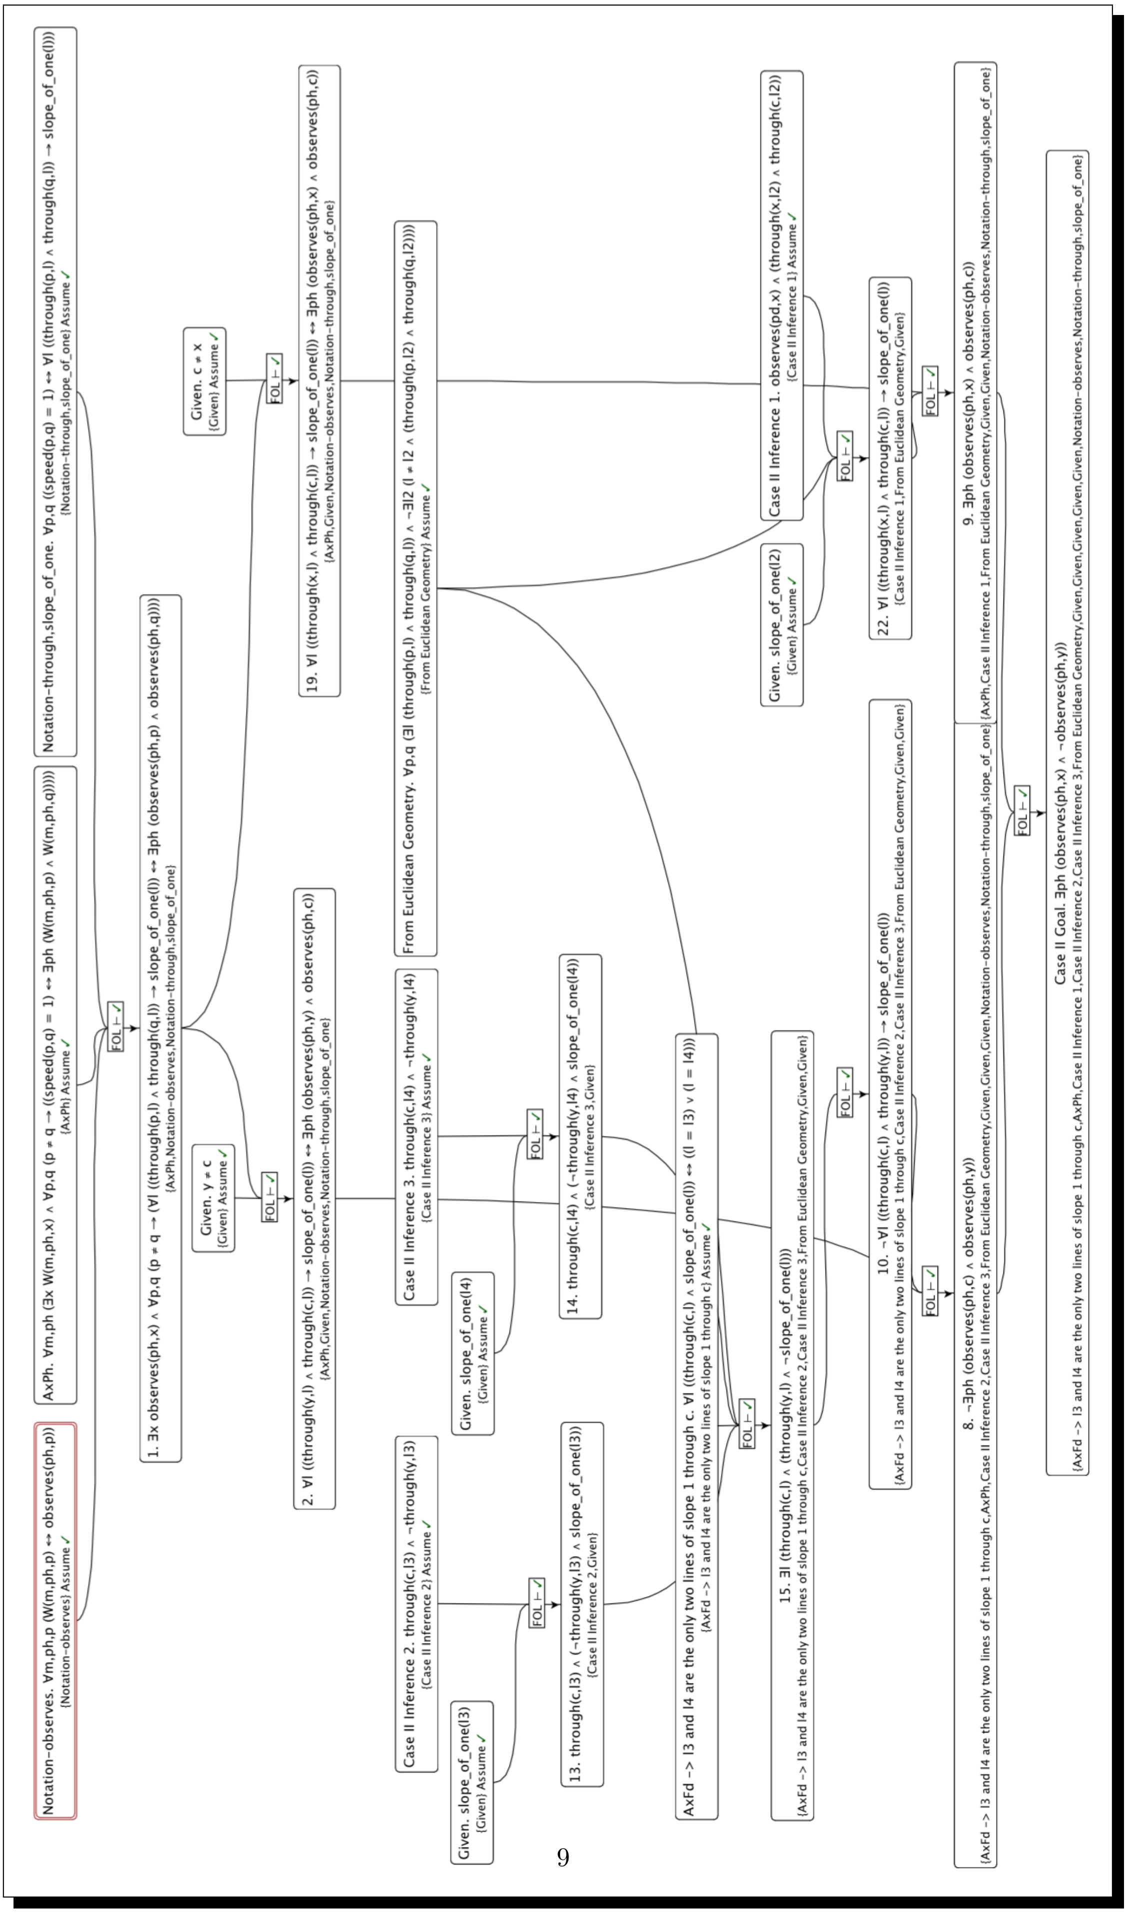

Figure 10: Proof for Case II

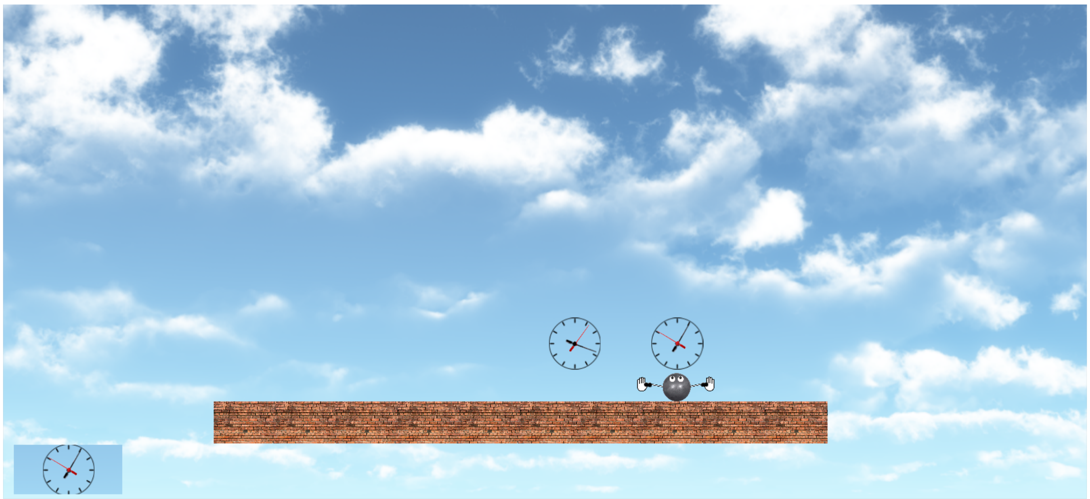

Figure 11: A demonstration of real-time reasoning with Vivid in PAGI World.

that if an observer observes a photon p at point x and x is on worldline l and y is not on worldline l, then that same photon p is not observable at y. Beginning with diagram Δ_0 (Figure 3), we apply the diagrammatic rule $[C_1]$ with the field axiom $AxFd$ to derive two new diagrams, Δ_1 and Δ_2, corresponding to the cases above, respectively. Formally, we have, **cases** from $AxFd : (\sigma_1, p_1) \rightarrow \Delta_1 \mid (\sigma_2, p_2) \rightarrow \Delta_2$, where y is either on a worldline or is not. These cases are clearly exhaustive, which satisfies the side condition necessary to implement $[C_1]$.

> **Case I** (Figure 4): From diagram Δ_1, we claim a photon p exists at point x and use the diagrammatic rule $[Diagram - Reiteration]$ to derive Δ_3 (Figure 6). We proceed by using the **observe** [7] rule to extract sentential information from the diagram. Specifically, we extract two observations: **observe** $observes(p, x) \wedge through(x, l_1) \wedge \neg through(y, l_1)$ and **observe** $observes(p, x) \wedge through(x, l_2) \wedge \neg through(y, l_2)$. Through application of the *not_on_same_worldline axiom* with both of the observations obtained from Δ_3, we conclude $\neg observes(p, y)$, that is, m does not observe photon p at y, but does observe p at x. The semi-automated proof of this is formalized in Figure 9, in the hypergraph-based automated theorem prover Slate. Then, the set of bodies observed by m at x is different than the set of bodies m observes at y, completing the first case.

> **Case II** (Figure 5): In the case of Δ_2, x and y lie on the same worldline l_1 depicted in the diagram. However, $AxFd$ specifies another line of slope 1 through x. With this in mind, we apply the diagrammatic rule $[\Delta; \Delta]$ with $AxFd$ to derive diagram Δ_4 (Figure 7). Now, claiming photon p' exists at point c on the new worldline l_2 that goes through x (where $c \neq x$), and utilizing the $[Diagram - Reiteration]$, derives the new diagram Δ_5 (Figure 8). We now **observe** $observes(p', x) \wedge through(x, l_2) \wedge \neg through(y, l_2)$. $observes(p', x)$ implies $observes(p', c)$ as x and c are both on l_2. However, $observes(p', c))$ implies $\neg observes(p', y)$ since $AxFd$ tells us that any line between c and y has slope $k \neq 1$. Then, photon p' has a speed different from the speed of light which is prohibited by the light axiom $AxPh$. Hence, m observes photon p' through x but not through y, the semi-automated proof of which is formalized in Figure 10, again in Slate. This gives us a body in the set of bodies that m observes at x but which m does not observe at y. **QED**

4 PAGI World as a Simulator for Vivid

Psychometric AI (PAI, pronounced "pie") is an approach that gauges progress in AI by increasing the performance of AI systems on all tests designed for human intelligence (Bringsjord & Schimanski 2003, Bringsjord 2011). Psychometric Artificial *General* Intelligence (PAGI, pronounced "pay-guy") is likewise devoted to engineering AI systems capable of performing demonstrably well on *all* established, validated

[7] This rule is used to extract sentential information from diagrams. This is similar to what we do when we just look at something.

tests of intelligence and mental ability, even those which the systems have never seen before (Bringsjord & Licato 2012). PAGI tests can include much more than those which would typically appear on written exams. For example, being placed in a room and asked to make use of the available objects to figure out how to break out of the room might be considered a test of PAGI.

PAGI World is a simulation environment written in the Unity game development system to easily showcase artificially intelligent agents performing cognitive tasks. PAGI World is equally accessible to nearly all existing cognitive systems, as it is cross-platform and allows relevant AI agents to communicate with the world through TCP/IP, thus enabling AI technology to be tested/simulated using virtually any programming language. PAGI World has already been used to demonstrate some significant cognitive phenomena (e.g. see Marton, Licato & Bringsjord forthcoming, Atkin, Licato & Bringsjord forthcoming).

While Vivid serves as an abstract, formal logico-mathematical framework for the diagrammatic reasoning we employ, PAGI World promises to serve as an environment in which to construct and manipulate particular diagrammatic content. This makes it ideal as a simulator for Vivid proofs. For example, Figure 11 pictures a task in PAGI World containing two clocks, which the artificial agent views. The agent is able to extract relevant states σ_i from the clocks, putting them into Vivid formulae. Using those formulae, the artificial agent is able to determine which of the clocks is showing a later time.

In particular, this PAGI World-Vivid connection offers the enticing possiblity of simulating special relativistic kinematics. With small on-screen velocities representing velocities close to that of light, the simulation would demonstrate the famous triad of special relativistic effects; that is, that moving clocks slow down, moving 'spaceships' shrink, and moving pairs of clocks lose their synchrony. Most importantly, the simulation, since it concretizes the *diagrammatic* representation of a sequence of steps in a Vivid proof, is automatically verifiable at each step. This approach to computational simulation explicitly represents a 'proof in motion' (recalling the Curry-Howard Isomorphism), and is thus at once verifiable and simulation-based.

5 Next Steps

Ultimately, our goal is exactly the rather ambitious one stated in (Govindarajalulu et al. 2014): engineer intelligent computing machines capable of not only verifying formal reasoning, but *discovering* proofs.[8] We believe that such discovery will be much easier to obtain when the diagrammatic information that is historically central to human creativity in the formal sciences is made susceptible to mechanical manipulation. Hitherto, this manipulation, sadly, has been, save for a few rare and rarefied exceptions, restricted to mere linguistic processing. To put the desired future another way, we are seeking, as a framework for formal science, a comprehensive, heterogeneous logic of the type that Leibniz, throughout his professional life, dreamed of (Bringsjord & Govindarajulu forthcoming).

This paper takes steps in that direction by: introducing heterogeneous deduction to a community that has been exclusively focused on linguistic/sentential formalisms, despite the stark reality of the role of diagrams in rigorous reasoning in physics; demonstrating the formally verifiable heterogenous proof of a theorem of special relativity; and introducing a graphical simulator for such proofs. As a domain-independant project, the implementation of an *automated* theorem prover for Vivid languages will of course be necessary to meet our goal. In addition, we will be seeking ways of merging the advantages of hypergraphical natural deduction, as embodied in the Slate system used for the RAIR Lab's earlier, first wave of work in logicist physics, with diagrammatic elements.

Visualizations of various special-relativistic effects should be feasible in PAGI World, such as the pole-and-barn paradox and the twin paradox. Simulating such effects has potential applications in education, gaming, and digital art. Once extended to general relativity, our simulation-by-proof approach promises a merger between the 'analytic' and 'numerical' traditions of relativity practice, with applications similar to that of the Lazarus Project (Baker, Campanelli & Lousto 2002), but with numerous advantages over it: verifiability; potentially increased accuracy; and graphical simulation via PAGI World.

[8] A full discussion, tied to the Four Color Theorem, now classically proved and verified, is provided in (Arkoudas & Bringsjord 2007).

References

Aiello, M., Pratt-Hartmann, I. & Benthem, J. V. (2007*a*), What is Spatial Logic?, *in* M. Aiello, I. Pratt-Hartmann & J. V. Benthem, eds, 'Handbook of Spatial Logics', Springer, Dordrecht, The Netherlands.

Aiello, M., Pratt-Hartmann, I. & Benthem, J. V., eds (2007*b*), *Handbook of Spatial Logics*, Springer, Dordrecht, The Netherlands.

Andréka, H., Madarász, J. X. & Németi, I. (2007), Logic of Space-Time and Relativity Theory, *in* 'Handbook of spatial logics', Springer, pp. 607–711.

Arkoudas, K. (2000), Denotational Proof Languages, PhD thesis, MIT, Department of Computer Science, Cambridge, USA.

Arkoudas, K. & Bringsjord, S. (2007), 'Computers, Justification, and Mathematical Knowledge', *Minds and Machines* **17**(2), 185–202.
URL: *http://kryten.mm.rpi.edu/ka_sb_proofs_offprint.pdf*

Arkoudas, K. & Bringsjord, S. (2009*a*), 'Vivid: A Framework for Combining Diagrammatic and Symbolic Reasoning', *Artificial Intelligence* **173**(15), 1367–1405.

Arkoudas, K. & Bringsjord, S. (2009*b*), 'Vivid: An AI Framework for Heterogeneous Problem Solving', *Artificial Intelligence* **173**(15), 1367–1405. The url http://kryten.mm.rpi.edu/vivid/vivid.pdf provides a preprint of the penultimate draft only. If for some reason it is not working, please contact either author directly by email.
URL: *http://kryten.mm.rpi.edu/vivid_030205.pdf*

Atkin, K., Licato, J. & Bringsjord, S. (forthcoming), Modeling Interoperability Between a Reflex and Reasoning System in a Physical Simulation Environment, *in* 'Proceedings of the 2015 Spring Simulation Multi-Conference'.

Baker, J., Campanelli, M. & Lousto, C. O. (2002), 'The Lazarus Project: A Pragmatic Approach to Binary Black Hole Evolutions', *Phys. Rev. D* **65**, 044001.
URL: *http://link.aps.org/doi/10.1103/PhysRevD.65.044001*

Barwise, J. & Etchemendy, J. (1995), Heterogeneous logic, *in* J. Glasgow, N. Narayanan & B. Chandrasekaran, eds, 'Diagrammatic Reasoning: Cognitive and Computational Perspectives', MIT Press, Cambridge, MA, pp. 211–234.

Bringsjord, S. (2011), 'Psychometric artificial intelligence', *Journal of Experimental and Theoretical Artificial Intelligence* **23**(3), 271–277.

Bringsjord, S., Govindarajulu, N., Ellis, S., McCarty, E. & Licato, J. (2014), 'Nuclear Deterrence and the Logic of Deliberative Mindreading', *Cognitive Systems Research* **28**, 20–43.
URL: *http://kryten.mm.rpi.edu/SB_NSG_SE_EM_JL_nuclear_mindreading_062313.pdf*

Bringsjord, S. & Govindarajulu, N. S. (forthcoming), Leibnizs Art of Infallibility, Watson, and the Philosophy, Theory, & Future of AI, *in* V. Müller, ed., 'Synthese Library', Springer.
URL: *http://kryten.mm.rpi.edu/SB_NSG_Watson_Leibniz_PT-AI_061414.pdf*

Bringsjord, S. & Licato, J. (2012), Psychometric Artificial General Intelligence: The Piaget-MacGyver Room, *in* P. Wang & B. Goertzel, eds, 'Theoretical Foundations of Artificial General Intelligence', Atlantis Press, 8, square des Bouleaux, 75019 Paris, France, pp. 25–48.
URL: *http://kryten.mm.rpi.edu/Bringsjord_Licato_PAGI_071512.pdf*

Bringsjord, S. & Schimanski, B. (2003), What is Artificial Intelligence? Psychometric AI as an Answer, *in* 'Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI)'.

Bringsjord, S., Taylor, J., Shilliday, A., Clark, M. & Arkoudas, K. (2008), Slate: An Argument-Centered Intelligent Assistant to Human Reasoners, *in* F. Grasso, N. Green, R. Kibble & C. Reed, eds, 'Proceedings of the 8th International Workshop on Computational Models of Natural Argument (CMNA 8)', Patras, Greece, pp. 1–10.
URL: *http://kryten.mm.rpi.edu/Bringsjord_etal_Slate_cmna_crc_061708.pdf*

Govindarajalulu, N., Bringsjord, S. & Taylor, J. (2014), 'Proof Verification and Proof Discovery for Relativity', *Synthese* pp. 1–18. Hard-copy version forthcoming; published in e-print form before release 18 Apr 2014.
URL: *http://link.springer.com/article/10.1007%2Fs11229-014-0424-3*

Licato, J., Atkin, K., Cusick, J., Marton, N., Bringsjord, S. & Sun, R. (Submitted for Initial Review), 'PAGI World: A Physically Realistic Simulation Environment for Developmental AI Systems', *Journal of Artificial Intelligence Research* .

Marton, N., Licato, J. & Bringsjord, S. (forthcoming), Creating and Reasoning Over Scene Descriptions in a Physically Realistic Simulation, *in* 'Proceedings of the 2015 Spring Simulation Multi-Conference'.

Sloman, A. (1971), Interactions Between Philosophy and Artificial Intelligence, *in* 'Proceedings from the Second International Joint Conference on Artificial Intelligence'.

Stannett, M. & Németi, I. (2014), 'Using Isabelle/HOL to Verify First-Order Relativity Theory', *Journal of automated reasoning* **52**(4), 361–378.