

Data Alchemy 201: Glass Bead Games

Victor Schetinger*
St. Pölten University of Applied Sciences

Rodrigo Oliveira de Oliveira†
Universidade Federal de Santa Maria



Figure 1: A random set of 100 motif beads being visualized in three different perspectives or scaffolds: Tarot (left), I Ching (center), and emoji (right). Each motif has a vector representation in an embedding space, which allows it to be “binned” into a different system of representation using the shortest cosine distance to other semiotic objects (that also have embedding vectors). On the left, a motif is assigned to a minor arcana (think of a deck of poker cards) with a rank from A, 2, 3, ..., 10, J(ack), Q(ueen), K(night), K(ing), and a suit: swords (or spades), cups (also hearts), pentacles (diamonds), and wands (clubs). In the center, the same motif would have a different representation in the space of the 64 hexagrams of the I Ching. In the right, instead we represent each motif as a string of its k -nearest (here $k=10$) emojis, starting from the center. Therefore, each bead is part of a complex network of semiotic relations that define its unique “identity” or “self”, seeded from its nominal description (i.e. the string with the motif’s name).

ABSTRACT

In this paper, we present an approach to explore data and latent spaces by playing games on them, inspired by Herman Hesse’s 1946 book “The Glass Bead Game” (GBG), which talks about a lineage of games that are played by making connections between ideas. We use vector embeddings to “glue” different datasets and provide semantic lattices on which games can be designed and played. Through a narrative exposition, we explain the fundamentals of our framework, showcase intermediary results, and discuss our long-term goals with this line of research. Fundamentally, we are still pursuing the same vision from the previous alt.vis papers, which is to look for meaningful semiotic structures within latent spaces, but in this time we propose a practical approach to get there based on visual analytics, games, and category theory. We provide reproducible instructions to bootstrap our approach on the Thompson Motif Index (a rich catalog of narrative motifs) and play card games with it using the Tarot minor arcana. In honor of the GBG, we frame them as “motif beads”, and show the countless possibilities that they allow.

1 INTRODUCTION

This paper can be thought of as the third (or fourth) in a series published on alt.vis that deals with digital semiotics, transdisciplinarity, and of course, visualization. Through playful exploration, we have been hypothesizing about how the latent spaces of models can be visualized in interesting ways and if they contain some forms of archetypal representations that could help us learn about

how humans construct meaning. This is still a very important line of research within the fields of visualization and machine learning [42, 45], but the truly interesting questions come from transdisciplinary contexts.

Recent work has shown that there might be some form of universal geometry in the topology of (large enough) language models [9, 24], and that through the lens of information thermodynamics such *cognitive topologies* might be inevitabilities even in the biological setting [1, 12, 19, 26, 31, 32]. Compression, mnemonics (we knew it! [39]) and memoization under material constraints seem to be at the root of all morphogenesis [17, 27]. Even consciousness seems to be limited by some bandwidth bottleneck [23, 51], which could also be expressed in this materiality [11]. What is even more interesting is that [the hunt for such “causal patterns”, abstract objects that can exist through different modes of representation, at different scales, and instantiated in various media, can be a research program in itself \[25\].](#)

How does one begin to seriously pursue this quest for virtual archetypes from a research perspective? In our previous papers, we developed the intuitive notions required to ground this semiotic investigation within visualization. In this paper, we are packaging a few different lines of in progress transdisciplinary studies within a ludic framework that, when successful, can allow one to carry out *actual research* and teaching by making and playing games. **Our core idea is to develop expressive semantic game design tools using vector embeddings that, with the help of [category theory](#), can be used to “glue” datasets of interest and latent spaces, allowing for the translation of visual analytics tasks into game domains and vice versa.** **PS.** [To aid the reader in understanding the concrete aspects of this work, and to more easily digest these complex ideas, we provide \[here an AI generated audio abstract \\(~6m\\)\]\(#\) and a demo showcasing our vis and some demo games \(~23m\).](#)

*e-mail: victor.schetinger@fhstp.ac.at

†e-mail: rooliveira@gmail.com




Figure 2: Schematics of the process of construction of the enriched Motif Beads described in Fig. 1. Every column can be thought of its own category, or in more practical terms just different versions of the dataset with more fields. On the left, we start with a dataset of the individual motif names (as strings) that are part of the Thompson Motif Index [49]. This string is always going to be the primary key of the motif, and therefore its highlighted in black and placed as the first field. First, we calculate the vector embedding for it using our preferred embedding model (green), then we can use any form of image generation to attach an image to this motif (blue), and in our language we get the “Motif Beads”, which can be thought of individual game assets at this point. On the right side we do a similar process with any dataset we want to use as symbolic system, in this case the Tarot Minor Arcana. We can then use the card that has the smallest cosine distance to our motif (in the vector embedding space), to attach a single rank and suit (in red, e.g., King of Cups, Two of Swords, etc) and make it playable in essentially any card game that uses a “regular” deck of playing cards. The same process can be repeated additively for any symbolic system, such as the I Ching or our emoji representation.

Knowledge Games

The field of game research (of which gamification is a subtopic) has many terms for different types of “games with a purpose”: educational games, serious games, engagement games, and so on, depending on the scope, goals and methodologies used. For our purposes, we are mostly interested in what Karen Schrier conceptualizes as “Knowledge Games” [41]: games with the potential of generating both internal knowledge for the player (through the development of skills and competencies, cultural acquisition) and external knowledge for society (through the generation of data and insights). It can be considered as adjacent to learning games and crowd-sourcing or citizen games, but taking a transdisciplinary stand these distinctions are not that important. **What is important is that in all these different scenarios, the key challenge is effective game design** [3, 28, 30, 41], from a theoretic perspective (knowing *how* to design and implement the key mechanics), but most seriously from an operational perspective (having the right resources to make, polish and maintain the games).

Our proposed approach aims to give all the power we can to game designers, reducing the cost of production without prescribing *how* they should do their games (or what they should look like). We are building some powerful tooling that, when used to make games, should naturally support the elevation of any game design into a knowledge game (in some domain). Many games already collect massive amounts of data, be they “serious” [2] or not, but using the paradigm of persistent interactions [33] we intend to create a bridge between visual analytics and game design without forcing too many constraints on either side, ultimately treating the game designer as an expert with the task of making **fun** games (that is non-negotiable). Furthermore, methodologies for “general purpose” exploration of emergent phenomena with visualization [13] could be easily leveraged within our framework. The remainder of the paper discusses our current progress in this direction.

2 BACKGROUND

To help the reader situate themselves, here we provide a quick recap of our previous alt.vis publications. For us, each paper captured a very interesting moment of our research lives and interests, as if they different snapshots or data points into the last 5 years of the

“General Purpose Transformer” revolution that is looking a bit too fast at the moment. We can look back at 2021 as simpler times where we were locked in our homes worrying if COVID would ever end, still unaware of what was about to come. Because of that, and because trilogies are cool, we decided to number our first paper as 0, so that the current one would be the III:

2021 alt.vis.0: Xenakis Just as the GBG mythical origins start with ludic music exercises, our first submission to alt.vis plays with the idea of “hearing a city” by transforming the visual structure of its streets into musical loops. We consider it a preamble to the current research from more simple (pre-LLM) times, but the seeds of our multimodal, trans(mutational)disciplinary explorations are clearly there. It also features some of the groundwork for another line of research into the “informational thermodynamics” of the human-in-the-middle, which we are still developing in the context of persistent interactions [33].

2022 alt.vis.I: Data Alchemy 101 When text-to-image models started to become accessible, we became obsessed with the idea of **visualizing the semiotic content crystallized within these models**, and experimented extensively with this using diffusion techniques. Visualizing latent spaces was (and still is) a challenge, and we brought forward the concept of “data alchemy”, which amounts to applying alchemical thinking to digital forms of representation to tackle this transdisciplinary challenge. As a largely theoretical paper, we produced different types of media to aid the digestion of our exposition: a paper, a short comic, and a rhizomatic miro board with extensive content, easter eggs and rabbit holes hidden inside it.

2023 alt.vis.II: n-Walks in the Fictional Woods ChatGPT was released to the public at the end of 2022, and this brought a shift in our focus from images to text. Still, we were interested in understanding what was inside of these models, what secret archetypal knowledge they had distilled from all our written media and encoded in the topology of their latent spaces. We used the semiotics of Umberto Eco [14] to continue our exploration into transdisciplinary waters, arguing that narratives are a powerful medium for compressing causal

knowledge and a possible origin for scientific thought. We also developed a prototype to visualize clusters of narratives generated by LLMs, which later allowed us to further analyze the storytelling capabilities of chatGPT through storyboards [40], and this motivated the solution proposed in the current paper that uses Folk-Literature motifs (more about this in the next sections).

In this “volume”, we chose the novel “The Glass Bead Game” (sometimes also known as *Magister Ludi*), written by Nobel laureate Herman Hesse and published in 1946. Hesse writes about an optimistic utopian future where humanity has attained global peace in the realization of its ultimate purpose: to play the Glass Bead Game (GBG). The game itself is tangentially described and contextualized in vivid expositions as it evolved through the ages, **but its rules and gameplay are never truly explained**. What is clear is that it is a game that synthesizes the whole of human culture and learning, and players compete in music-like duets by weaving connections between ideas using a **visual language of representation**.

3 THE GLASS BEAD GAME

In the book’s introductory chapter (*The Glass Bead Game: A General Introduction to Its History for the Layman*), we learn that the GBG was invented sometime around the early XX century by a music teacher in Cologne, Germany, as a memory-and-improvisation exercise for students. Colored glass beads would represent different kinds of musical motifs and notation, which would be picked and arranged by a player (e.g. the teacher), and the other player (e.g. a student) would have to respond by continuing the piece in a sort of back-and-forth musical dialogue of themes. A decade or two later, likely around the 1910s in the book’s timeline, mathematicians picked up interest in the game and elevated it into complex exercises of logics and pattern recognition, exchanging the symbolic meanings of the beads with their own notations, and adapting the rules. Soon, it spread to many domains of human knowledge so that a myriad of variants existed for nearly all scholarly disciplines.

The Game of Games

Over the next centuries, we are told by the narrator, the game has always existed in one shape or another, sometimes enjoying mass popularity, sometimes waning into cliqued communities, but eventually it was re-unified as a sort of **universal symbolic medium**, with its own lexicon of ideograms and **visual grammar** that attempted to capture the whole of human culture, and although the name “*Glass Bead Game*” was still used to refer to it, it did not involve glass beads anymore at all. One of the most elucidating passages that can help the reader form a picture of the game is from the same introductory chapter:

Under the shifting hegemony of now this, now that science or art, the Game of games, had developed into a kind of universal language through which the players could express values and set these in relation to one another. Throughout its history the Game was closely allied with music, and usually proceeded according to musical or mathematical rules. One theme, two themes, or three themes were stated, elaborated, varied, and underwent a development quite similar to that of the theme in a Bach fugue or a concerto movement. A Game, for example, might start from a given astronomical configuration, or from the actual theme of a Bach fugue, or from a sentence out of Leibniz or the Upanishads, and from this theme, depending on the intentions and talents of the player, it could either further explore and elaborate the initial motif or else enrich its expressiveness by allusions to kindred concepts. [22]

In Hesse’s philosophy, the game is a way for humans to decode wisdom within culture by externalizing it collectively and *playing* with it. It has a fundamentally empirical and inexhaustible quality that requires engaging with peers and our cultural latent spaces in search for *Geist*. This resembles a romanticized, idealized view of how some areas of vis see themselves, where expert users use their tacit knowledge to grind insights from data. In any case, before the Glass Bead Game arrived at this stage of Hegelian engine, it slowly evolved with humanity while bootstrapping all its necessary components. Let’s pretend we are trying to help this bootstrapping by exploring the state-of-the-art of different disciplines and testing what kinds of Glass Bead Games are currently possible.

Data Alchemy 201

The *framing* of the GBG we are using for this paper developed organically from experiments in game design using LLMs and vector embeddings to understand how a “semantic” game engine could look like in practice (in addition to the graphics engine, the physics engine and other components of the architecture of modern game engines), and what sorts of novel games it would allow. This *frame* turned out to be a very opportunistic way to deal with the challenges of exploring latent spaces, of which three of them (and our proposed solutions) are central to this paper:

High Dimensionality All internal forms of representation used by modern models (of which vector embeddings are the most common, and our focus) are in the order of hundreds of dimensions, and there are currently no methods to visualize and compare their semiotic content (i.e. what symbolic, mnemonic associations it is making) in a satisfactory way. Dimensionality reduction methods such as t-SNE can be very useful for “glimpsing” at things, but will always collapse complex relational topologies that are of interest to us, and therefore can provide only “thin slices” of information. **Take-away:** there are no good solutions yet and they will likely involve highly interactive exploratory techniques. **We think visual analytics can help us here.** [10, 33, 48]

Ephemeral Results Every week a large amount of new models and Machine Learning papers are released, and many become obsolete before the community even has time to digest them. The idea of a state-of-the-art that has any form of consensus is currently a logistical impossibility in many fields, as if they were constantly being subjected to distributed denial of server (DDoS) attacks. Any solution that intends to provide a useful, general perspective on the inner workings of latent spaces must be grounded outside of this Heraclitian flux. **Take-away:** there is a need for research methodologies in this area that allow researchers to work at their own pace and investigate phenomena that can be generalizable, rather than chasing the latest thing hoping to be relevant. **We think that Category Theory can help us here.** [6, 18, 36]

Human-in-the-Loop One of the main takeaways from the whole of the visualization literature is the importance of humans actively participating in tasks, processes, decisions, etc. As our average attention spans dwindle and we conveniently delegate more to automated systems, it becomes ever harder to escape the positive-feedback loops of models teaching models to regulate other models that synthesize data to train models, and so on, to simulate some human dimension. It is essential to find ways to engage people to participate in the processes that shape and align these models, especially when they are *playing* with our psychology and precious cultural substrates. **Take-away:** We need humans engaged in alignment tasks that are cognitively

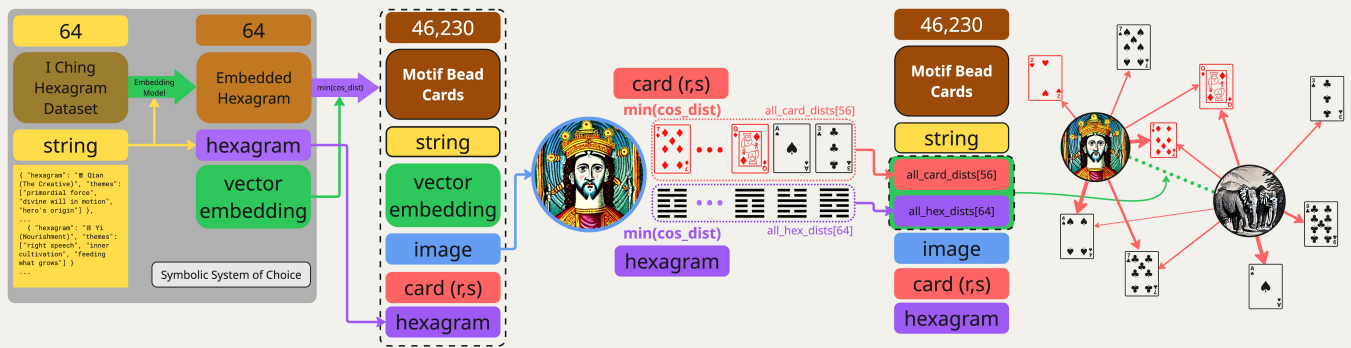


Figure 3: (left) The process of adding another symbolic system (here the I Ching) to a motif bead is the same as in Fig. 2. (center) While assigning a specific card with a rank and suit (or hexagram) to a motif amounts to picking the element from each set with the minimum distance to it, one could also keep the whole ordered array of the closest cards (or hexagrams, emojis, etc.), and use it as a proxy for the embedding vector, dropping it altogether. This approach can work as a form of dimensionality reduction going from the hundreds of dimensions of the vector embedding to the ordered set of 56 cards or distances. (right) The sets of distances can be used to establish relationships between motifs that will still carry semantic meaning, and can therefore provide for the design of unique game mechanics in their own right while reducing memory usage significantly.

relevant and rewarding, while at the same time useful. **We think that game design can help us here [41,43,44,47].**

This work can be thought of as a “response” to the questions posed by **Data Alchemy 101** [39], not in a “here are the answers we found” way, but as praxis (we are still looking for the answers, by the way). There are many devices at work within Hesse’s book that give it a certain **hyperstitional vitality**, but in particular the constant implication that *the game was already evolving in our times*, and that it could take practically any shape. Maybe some existing boardgames such as Concept, Dixit, Codenames, or even Scrabble could feature as distant ancestors in the utopian GBG family tree in our timeline? Why not make your own? How would a GBG for biologists look like? For lawyers? What sort of GBG can you conceive that is fun to you? **From a teaching perspective, this is a very rich source of exploration that can help students develop many transdisciplinary skills.** If Data Alchemy 101 covers the theoretical basis for transmuting digital materials, in 201 we propose practical experimentation by conceptualizing, making, playing and analyzing GBGs.

4 FROM SEMANTIC ENGINES TO PLATONIC SPACES

The feasibility of using LLMs for game development is still quite limited due to (1) the intense use of computational resources when running locally and (2) the scaling costs of depending on API calls. When exploring the idea of semantic engines for game development, this was an important issue to consider. Our solution was to focus on the use of **vector embeddings** (Fig. 2), because while a small transformer model with token prediction will be on the order of a few GBs with questionable quality, vector embedding models (that is, models that only take a string and output a vector embedding without the word prediction part) can be a few dozen MBs. Furthermore, vector embeddings can be calculated offline, and because operating with them is just simple vector arithmetic, real-time inference can be avoided in many cases, which is *very* convenient for game development.

So, what can you *actually* do with vector embeddings within game design? The most straightforward way to use them is to compare the distance of something with something else from a semantic perspective and use this as a trigger for some behavior. For instance, when generating a new item, one could check $\text{dist}(\text{item.name}, \text{"light"}) < \text{dist}(\text{item.name}, \text{"darkness"})$ to decide if the attribute of “light” or “darkness” will be added it and then have some game mechanic associated with this attribute. Or, to dynamically determine if a certain creature could spawn in a biome,

one could test, for example, if $\text{dist}(\text{creature.name}, \text{"marsh"}) < 0.15$. This is already very useful for procedural generation of content [46] and mechanics, but it might not justify having to deal with the extra complexity it adds to the project if there is not some sort of *ontology* on top of it supporting the game design process. The price we pay for ditching the rest of the model machinery (transformer or otherwise) after the token-to-embedding part is the loss of higher-level articulation.

Scaffolding Meaning

The first challenge we encountered in our experiments was to elevate vector embeddings to a working language for game design without adding too much complexity. Say you have some crafting mechanic that allows arbitrary combinations of two items with embeddings A and B , something that from a vector embedding perspective can be straightforward to achieve (just take the average of the embeddings as $(A+B)/2$ and normalize it). Now you have coordinates C in a high-dimensional space, but what do you do with them if you don’t have the machine to turn them into a string? If in the context of your game there is an object in the latent space between these two coordinates, you could select it as the resulting combination, but otherwise there is not much you can do. You need to have “something” there, an object D so that $\text{dist}(C, D) < \text{dist}(C, A)$ or $\text{dist}(C, D) < \text{dist}(C, B)$ to provide a meaningful combination mechanic. A solution for this problem is to have lattices that “scaffold” the space in such a way that there will *always* be an object D .

For example, one could have a whole English dictionary with embeddings loaded in the memory (e.g. **GloVe**) so that if “fire” and “water” were combined, then “steam” might pulled out as the result. The issue with this approach is that not only we are not guaranteed to have good results, but trying to meaningfully cover the whole space would quickly bring us back to having to keep GBs of data in the memory. Granted, it still does not need real-time inference or GPU memory, but this limits the feasibility of this solution. For a dictionary-like lattice, one could take advantage of **Zipf’s Law** and greatly reduce the amount of used words (e.g., top 10% most used words), which can be a good compromise for tackling the sparsity issue of the semantic engine. Still, a dictionary is not the only set that could be used as a lattice, and we shall look into other interesting options that can be combined with it.

The Arrows of the Beholder

What we call a “Symbolic System” of choice in Fig. 2 and 3 are just different lattices one could use to scaffold the latent space, and

conversely Fig. 1 shows how this choice can give us different views over the same data. This is where the beauty of [Category Theory](#) (CT) comes in. While a non-superficial exposition of CT is far beyond the scope of this paper (see our Appendix B for a taste), it can be thought of as a sort of “mathematical relativity” (as in Einstein’s relativity). If classical Cartesian mathematics takes the 0-dimensional point as its starting object, CT says that the point is just a “hole” left by an arrow pricking (pointing) there, and the arrow is the real important object [29]. In a way, this arrow is just a regular *(target, source)* edge on a directed graph, but, in the hands of category theorists it becomes a magic wand and takes a life of its own, being able to create powerful abstractions and bridges between different parts of mathematics. In a way, through the power of the arrow *almost* all mathematics can live within CT.

Vickers et al. did a wonderful job contextualizing CT for vis [50], but maybe they were a bit ahead of their time, because somewhat surprisingly CT is to this day underrepresented within vis. We say surprisingly because it is deeply intertwined with graphs and graph theory (every category has an underlying directed graph representation!), one of the most important subjects of vis. This is a synergy that we plan to leverage to engage the community. More importantly, within the heart of CT is the idea that *some-thing is equal to the sum of all different ways of seeing it* (i.e. all paths of arrows coming and going through it, formalized by the famous [Yoneda lemma](#)). Therefore, the more we come up with different ways of seeing something, the more we approach its *true being*. If, as recent research has indicated [24], there are universal, *platonic*-like forms that emerge in latent spaces, then modeling new, different perspectives (through scaffoldings of symbolic systems or other ingenious methods we can come up with [15]) can be a methodology to look for them. Category theory gives us powerful tools to formalize these notions [35], such as [categorical limits](#), and to build bridges [8] to other physical and biological systems outside of latent spaces [25].

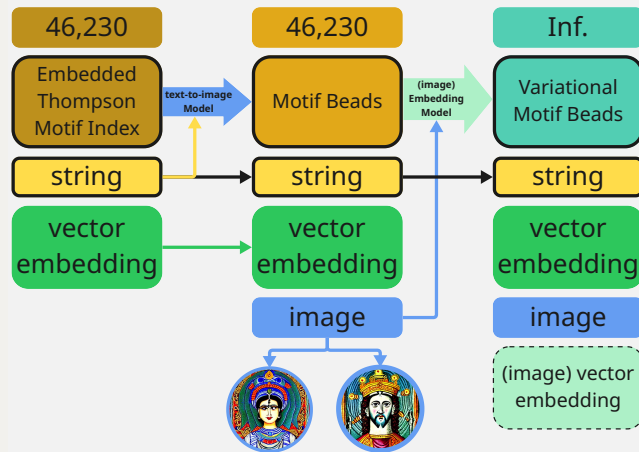


Figure 4: Because the same motif can have infinite variations of images when run it as a prompt for a text-to-image model, this can be exploited to create another layer of emergent game mechanics. Using the generated image as input for a multi-modal model (such as CLIP [34]) that can produce its own embeddings, and attaching it back to the bead, we can have instances of the same motif that have different behavior. Moreover, it is expected that the behavior is still somewhat close semantically to the original motif, but also dependent on the visual content of the image.

5 PLAYING WITH INDRA’S PEARLS

At this point in our exposition we would like to be able to [import from](#) our previous alt.vis papers, especially *n-Walks in*

[the Fictional Woods](#) [38], where we discuss the role of stories, myths and narratives in the development of human knowledge. Both the Tarot and the [I Ching](#) (the Chinese Book of Changes) also have a special place in our series, as they can be considered at the center of the metaphysical and oracular traditions of the West and East, respectively. Since the early days of mankind there has been a deep connection between playing games, telling stories and performing oracles [20], and the I Ching is the oldest form of oracle still in use. The Tarot is much more modern (and even more so its oracular use), but nevertheless it encodes many layers of European culture and myth [16] and can be easily co-opted for our game design ends – it is a deck of *playing cards*, after all!

In our current approach, we have built a few different scaffolds with different granularities over the latent space. The Tarot with 56 cards and the I Ching with 64 hexagrams are relatively *low-dimensional*, and we would like to treat them as possible *basis of representation* (Fig. 3), although achieving some sort of linear independence with them is still a challenge. What might not have been evident for the reader until now is that even a dataset that is intended to be studied (as the object of a knowledge game, for example) will also be a sort of symbolic system or scaffolding. In the spirit of Category Theory, the distinction between object and subject is purely utilitarian. In any case, we could say that the “subject” in this exposition is our dataset of *motif beads*, built from [Thompson’s Motif Index \(of Folk-Literature, or TMI\)](#) [49], which contains an exhaustive list (46,230 in the version used, the top number in Fig. 2) of motifs found in folk tales around the world.

Motif Beads and the Building Blocks of Stories

One of our goals with this project is to release sets of tools (both theoretical and technical) to assist in the development of knowledge games using embedding vectors. Motif beads (and consequently Motif Bead Cards, as in Fig. 2 and Fig. 3) are just a nice way that we chose to implement some of our ideas and package them within a GBG framework (and *fluff*). In Appendix A we provide instructions how to build them from scratch, so that people can experiment with them. Besides being interesting datasets for visualization on their own right (Fig. 1), we think they can make great assets for game development representing all sorts of entities such as items, spells, events, but really anything is possible, just think of what kinds of things cards can represent in card games. Also, one is not prescribed to use *all* motifs, or *any at all*, and for certain games and settings there might be subsets of a few hundred or a few dozen that can support everything in the game design. They can also be completely hidden from the player and used under the hood to power procedural generation and other kinds of mechanics, and even have multiple variations with different behavior (Fig. 4).

The TMI can be thought of as the diametrically opposing approach to narratology in relation to the famous “Monomyth” (or Hero’s journey) [7], as it goes for maximum granularity in the characterization of stories. Its almost as if the TMI is the adjoint functor of the Monomyth, in Category Theory parlance (i.e. one classification refines another via a pair of adjoint functors). This makes it a specially interesting candidate for our exploration, as game designers commonly study the Monomyth as part of their foundations in storytelling. Therefore, working (and playing, and trying to make games) with TMI data is helping them develop foundational knowledge! In addition, because the TMI is used by academics in different fields connected to folk studies and anthropology, this gives us a convenient transdisciplinary bridge, which is very much in the spirit of the GBG.

Wood is Below, Water Above

The main storyline in the GBG book is set in the pedagogical province of Castalia, run by a sort of monastic order of intellectuals and educators. It follows the life of Magister Ludi Joseph

Knecht as a retrospective biographical account from the utopian future, where he is considered a central figure that revolutionized the GBG. The Magister Ludi (literally, game master) is the highest authority in the GBG, and we accompany Knecht's *Bildungsweg* from a timid child through his ascension in the ranks of the order. Among his contributions to the GBG, Knecht introduced I Ching and its philosophy into the game after having studied it for many years under the tutelage of a master. In his early years as a student of the GBG, he had an epiphany where he saw every play performed in the game as a dialectic movement mirroring the changes in the I Ching:

I suddenly realized that in the language, or at any rate in the spirit of the Glass Bead Game, everything actually was all-meaningful, that every symbol and combination of symbols led not hither and yon, not to single examples, experiments, and proofs, but into the center, the mystery and innermost heart of the world, into primal knowledge. Every transition from major to minor in a sonata, every transformation of a myth or a religious cult, every classical or artistic formulation was, I realized in that flashing moment, if seen with a truly meditative mind, nothing but a direct route into the interior of the cosmic mystery, where in the alternation between inhaling and exhaling, between heaven and earth, between Yin and Yang, holiness is forever being created. [22]

This quote captures somewhat poetically the idea that we are bringing forth in this paper: people making and playing games, recording them, analyzing them and proposing new scaffoldings or perspectives can be part of an iterative process (almost a [Monte Carlo experiment](#)) towards universal knowledge constructions. We invoked the GBG because it felt too close to what we were doing and, when we looked into it more carefully, suddenly we were drawn to many unexpected interesting connections and resources that enriched our research. We hope that our exposition of this work in progress has been able to give the reader a somewhat clear picture of our ideas and ambitions. A full report of our methods and results should be forthcoming as publications as soon as our time and (human) resources allow.

ACKNOWLEDGMENTS

We are thankful for the space alt.vis provides for sharing this work even in its incomplete form, as it has been personally very challenging for us to keep up with the current pace of developments. Because we chose to work in a transdisciplinary setting, it makes it all the more difficult for us to be academically rigorous and up-to-date in many fields. Different AI models were used to assist in the writing of this paper: Gemini (NotebookLM) was used for literature review, chatGPT was used for discussion of ideas and LaTeX formatting, and various *vibe coding* platforms for making the prototypes.

REFERENCES

- [1] B. AlKhamissi, G. Tuckute, A. Bosselut, and M. Schrimpf. The LLM Language Network: A Neuroscientific Approach for Identifying Causally Task-Relevant Units, Nov. 2024. arXiv:2411.02280. doi: 10.48550/arXiv.2411.02280
- [2] C. Alonso-Fernandez, A. Calvo-Morata, M. Freire, I. Martinez-Ortiz, and B. Fernández-Manjón. Applications of data science to game learning analytics data: A systematic literature review. *Computers & Education*, 141, June 2019. doi: 10.1016/j.compedu.2019.103612
- [3] J. Alvarez and D. Djaouti. An introduction to Serious game definitions and concepts. *Serious Games & Simulation for Risks Management*, Dec. 2011.
- [4] A. Ayzenberg, T. Gebhart, G. Magai, and G. Solomadin. Sheaf theory: from deep geometry to deep learning, Feb. 2025. arXiv:2502.15476 [math]. doi: 10.48550/arXiv.2502.15476

- [5] G. M. Bergman. On diagram-chasing in double complexes, Feb. 2012. arXiv:1108.0958 [math]. doi: 10.48550/arXiv.1108.0958
- [6] F. P. Z. Camacho. Embedding ontologies using category theory semantics. Masters of science thesis, King Abdullah University of Science and Technology, Thuwal, Kingdom of Saudi Arabia, March 2022.
- [7] J. Campbell. *The Hero with a Thousand Faces*. Bollingen series. New World Library, 2008.
- [8] O. Caramello et al. The theory of topos-theoretic 'bridges': a conceptual introduction. *Glass Bead Journal*, pp. 1–14, 2016.
- [9] B. Chughtai, L. Chan, and N. Nanda. A Toy Model of Universality: Reverse Engineering How Networks Learn Group Operations, May 2023. arXiv:2302.03025 [cs]. doi: 10.48550/arXiv.2302.03025
- [10] L. Cibulski and S. Bruckner. Towards Understanding Decision Problems As a Goal of Visualization Design, July 2025. arXiv:2507.18428 [cs]. doi: 10.48550/arXiv.2507.18428
- [11] M. DeLanda and R. Braidotti. *Materialist Phenomenology: A Philosophy of Perception*. Theory in the New Humanities. Bloomsbury Academic, 2021.
- [12] B. DeMoss, S. Sapora, J. Foerster, N. Hawes, and I. Posner. The Complexity Dynamics of Grokking, Dec. 2024. arXiv:2412.09810. doi: 10.48550/arXiv.2412.09810
- [13] A. Diehl, A. A. Rahman, B. Bach, M. El-Assady, M. Kraus, R. S. Laramée, D. A. Keim, and M. Chen. An Analysis of the Interplay and Mutual Benefits of Grounded Theory and Visualization. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–18, 2024. doi: 10.1109/TVCG.2024.3452985
- [14] U. Eco. *Six Walks in the Fictional Woods*. Harvard University Press, 1994.
- [15] A. Ehresmann and M. Béjean. The glass bead game revisited: Weaving emergent dynamics with the mes methodology. *Glass Bead Journal*, (1), 2016.
- [16] H. Farley. *A cultural history of tarot: from entertainment to esotericism*. I. B. Tauris, London, 2009.
- [17] C. Fields, A. Goldstein, and L. Sandved-Smith. Making the Thermodynamic Cost of Active Inference Explicit. *Entropy*, 26(8):622, Aug. 2024. Number: 8 Publisher: Multidisciplinary Digital Publishing Institute. doi: 10.3390/e26080622
- [18] T. Gebhart, J. Hansen, and P. Schrater. Knowledge Sheaves: A Sheaf-Theoretic Framework for Knowledge Graph Embedding. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, pp. 9094–9116. PMLR, Apr. 2023. ISSN: 2640-3498.
- [19] M. Genkin, K. V. Shenoy, C. Chandrasekaran, and T. A. Engel. The dynamics and geometry of choice in the premotor cortex. *Nature*, pp. 1–9, June 2025. Publisher: Nature Publishing Group. doi: 10.1038/s41586-025-09199-1
- [20] J. L. Gillin and J. Huizinga. Homo Ludens: A Study of the Play-Element in Culture. *American Sociological Review*, 16(2):274, Apr. 1951. doi: 10.2307/2087716
- [21] J. Hagedorn. trilogy.
- [22] H. Hesse. *The Glass Bead Game*. Bantam Books, New York, N.Y., 1970. Traduzido do alemão *Das Glasperlenspiel* (1943); a tradução para o inglês é © 1969 por Holt, Rinehart and Winston, Inc. Esta edição Bantam foi publicada em acordo com Holt, Rinehart and Winston, Inc. Inclui prefácio de Theodore Ziolkowski.
- [23] D. D. Hoffman, C. Prakash, and R. Prentner. Fusions of Consciousness. *Entropy*, 25(1):129, Jan. 2023. doi: 10.3390/e25010129
- [24] R. Jha, C. Zhang, V. Shmatikov, and J. X. Morris. Harnessing the universal geometry of embeddings, 2025.
- [25] M. Levin. Ingressing minds: Causal patterns beyond genetics and environment in natural, synthetic, and hybrid embodiments, Feb 2025. doi: 10.31234/osf.io/5g2xj_v1
- [26] M. Levin. The Multiscale Wisdom of the Body: Collective Intelligence as a Tractable Interface for Next-Generation Biomedicine. *BioEssays*, 47(3):e202400196, 2025. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/bies.202400196. doi: 10.1002/bies.202400196
- [27] Z. Liu, Y. Liu, J. Gore, and M. Tegmark. Neural Thermodynamic Laws for Large Language Model Training, May 2025. arXiv:2505.10559 [cs]. doi: 10.48550/arXiv.2505.10559
- [28] T. Marsh. Serious games continuum: Between games for purpose

- and experiential environments for purpose. *Entertainment Computing*, 2(2):61–68, Jan. 2011. Publisher: Elsevier BV. doi: 10.1016/j.entcom.2010.12.004
- [29] G. Mazzola. Melting glass beads—the multiverse game of gestures and strings. *Glass Bead, Site 0: Castalia, the Game of Ends and Means*, 2016.
- [30] J. A. Miller and S. Cooper. Barriers to Expertise in Citizen Science Games. In *CHI Conference on Human Factors in Computing Systems*, pp. 1–25. ACM, New Orleans LA USA, Apr. 2022. doi: 10.1145/3491102.3517541
- [31] B. R. Munn, E. J. Müller, I. Favre-Bulle, E. Scott, J. T. Lizier, M. Breakspear, and J. M. Shine. Multiscale organization of neuronal activity unifies scale-dependent theories of brain function. *Cell*, 187(25):7303–7313.e15, Dec. 2024. doi: 10.1016/j.cell.2024.10.004
- [32] E. Murphy. ROSE: A Universal Neural Grammar. *Cognitive Neuroscience*, pp. 1–32, July 2025. Publisher: Informa UK Limited. doi: 10.1080/17588928.2025.2523875
- [33] I. Pérez-Messina, D. Ceneda, V. Schetinger, and S. Miksch. Persistent interaction: A conceptualization of user-generated artefacts in Visual Analytics. *Computers & Graphics*, 129:104232, June 2025. doi: 10.1016/j.cag.2025.104232
- [34] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever. Learning Transferable Visual Models From Natural Language Supervision, Feb. 2021. arXiv:2103.00020 [cs]. doi: 10.48550/arXiv.2103.00020
- [35] M. Robinson. Sheaf and duality methods for analyzing multi-model systems, Nov. 2016. arXiv:1604.04647 [math]. doi: 10.48550/arXiv.1604.04647
- [36] D. Rosiak. *Towards a Classification of Continuity and On the Emergence of Generality*. PhD thesis, DePaul University, Chicago, IL, December 2019. PhD Dissertation, Department of Philosophy, College of Liberal Arts and Social Sciences.
- [37] D. Rosiak. Sheaf Theory through Examples (Abridged Version), Dec. 2020. arXiv:2012.08669 [math]. doi: 10.48550/arXiv.2012.08669
- [38] V. Schetinger, S. D. Bartolomeo, E. S. d. Lima, C. Meinecke, and R. Rosa. \$n\$ Walks in the Fictional Woods, Aug. 2023. arXiv:2308.06266 [cs]. doi: 10.48550/arXiv.2308.06266
- [39] V. Schetinger, V. Filipov, I. Pérez-Messina, E. Smith, and R. O. de Oliveira. I Learn to Diffuse, or Data Alchemy 101: a Mnemonic Manifesto, Oct. 2022. arXiv:2208.03998 [cs].
- [40] V. Schetinger, D. Reis Pedrosa da Silva, S. Di Bartolomeo, E. Soares de Lima, C. Meinecke, and R. Rosa. Macunaima, papagaio ia, resolve crimes em praga. *Revista Geminis*, 14(3):21–37, 2023. doi: 10.14244/2179-1465.RG.2023v14i3p21-37
- [41] K. Schrier. *Knowledge games: how playing games can solve problems, create insight, and make change*. Tech.edu. Johns Hopkins University Press, Baltimore, 2016.
- [42] R. Sevastjanova. *Interactive Visual Investigation of Word Embedding Contextualizations in Large Language Models*. PhD thesis, Universität Konstanz, Konstanz, Germany, 2025.
- [43] R. Sevastjanova, H. Hauptmann, S. Deterding, and M. El-Assady. Personalized Language Model Selection Through Gamified Elicitation of Contrastive Concept Preferences. *IEEE Transactions on Visualization and Computer Graphics*, 30(8):5449–5465, Aug. 2024. Publisher: Institute of Electrical and Electronics Engineers (IEEE). doi: 10.1109/tvcg.2023.3296905
- [44] R. Sevastjanova, H. Schäfer, J. Bernard, D. Keim, and M. El-Assady. Shall we play? – Extending the Visual Analytics Design Space through Gameful Design Concepts. In *2019 IEEE Workshop on Machine Learning from User Interaction for Visualization and Analytics (MLUI)*, pp. 1–9. IEEE, Vancouver, BC, Canada, Oct. 2019. doi: 10.1109/mlui52769.2019.10075563
- [45] C. Shani, D. Jurafsky, Y. LeCun, and R. Shwartz-Ziv. From Tokens to Thoughts: How LLMs and Humans Trade Compression for Meaning, May 2025. arXiv:2505.17117 [cs]. doi: 10.48550/arXiv.2505.17117
- [46] D. F. Silva, R. P. Torchelsen, and M. S. Aguiar. Procedural game level generation with GANs: potential, weaknesses, and unresolved challenges in the literature. *Multimedia Tools and Applications*, Jan. 2025. doi: 10.1007/s11042-025-20612-9
- [47] C. Stoiber, M. Boucher, M. Keck, L. Amabili, R. G. Raidou, V. Filipov, V. Oliveira, V. Schetinger, and W. Aigner. EuroVis Workshop on Visualization Play, Games, and Activities 2025: Frontmatter. In C. Stoiber, M. Boucher, M. Keck, L. Amabili, R. G. Raidou, V. Filipov, V. Oliveira, V. Schetinger, and W. Aigner, eds., *EuroVis Workshop on Visualization Play, Games, and Activities*. The Eurographics Association, 2025. doi: 10.2312/visgames.20252015
- [48] D. Streeb, M. El-Assady, D. A. Keim, and M. Chen. Why Visualize? Untangling a Large Network of Arguments. *IEEE Transactions on Visualization and Computer Graphics*, 27(3):2220–2236, Mar. 2021. doi: 10.1109/TVCG.2019.2940026
- [49] S. Thompson. *Motif-index of folk-literature; a classification of narrative elements in folktales, ballads, myths, fables, mediaeval romances, exempla, fabliaux, jest-books, and local legends*. Indiana University Press, Bloomington, 1955.
- [50] P. Vickers, J. Faith, and N. Rossiter. Understanding Visualization: A Formal Approach using Category Theory and Semiotics. *IEEE Transactions on Visualization and Computer Graphics*, 19(6):1048–1061, June 2013. arXiv:1311.4376 [cs]. doi: 10.1109/TVCG.2012.294
- [51] J. Zheng and M. Meister. The unbearable slowness of being: Why do we live at 10 bits/s? *Neuron*, p. S0896627324008080, Dec. 2024. doi: 10.1016/j.neuron.2024.11.008

A PARTIAL RESULTS, TECHNICAL DISCUSSION AND DIY

We have many interesting preliminary results, including the visualization seen in Fig. 1, many scripts to explore the datasets, and MVP games. A video should be included with this paper to showcase them. We intend to opensource everything ASAP, but it needs to be cleaned up and organized because this project was heavily *vibe coded*, but the alternative was not to make things and we did not want to go that way. Fortunately, our approach is intended to be easily reproducible using similar sources for the datasets, and here we provide instructions to do so. The results obtained in [24] suggest that the choice of embedding model or specific dataset used for knowledge games should not matter in the *long run* towards finding interesting semiotic structures in latent spaces, and in developing transferable skills. The use of tools from category theory should allow us to treat every different implementation and variation as some local measurements that can be (at least partially) generalizable towards global structures [4].

All that being said, consistency is always good and helpful when it does not become too constraining. The most important thing to guarantee usable results is to use the same embedding model. In our case `text-embedding-nomic-embed-text-v1.5@q8_0` (146mb) was used, and we want to text the effect of switching to lighter versions soon. For the datasets, here is a breakdown:

- **Thompson Motif Index** [49]: this repository [21] contains the whole TMI in a few different formats, plus some other folk studies resources that can be useful. In our case, **we used the motif names (with their indexing) as input to get the vector embeddings**, and added them as an extra column for simplicity;
- **Tarot**: we chose this dataset for simplicity, but it should be noted that it uses the **Raider-Waite** and this has implications for its performance as a symbolic system. It is not our favorite version of the Tarot and it has its own idiosyncrasies (e.g., cards XI and VIII are switched). **The whole json entry for each card was used to get the vector embedding**. It should also be noted that for our visualization we switched the displayed cards for a public domain version of the Marseilles tarot for visual simplicity (Raider-Waite contains complex scenes for each minor arcana card);
- **I Ching**: There are many versions of the I Ching but the one most used in the West is likely Richard Wilhelm’s translation (with a foreword by Jung). It has entered the public domain in 2020 and this repo conveniently packs the text of each hexagram in a json. **As input for the vector embeddings, we used the whole json entry for each hexagram, including its different representations and changing lines**.
- **Emojis**: for our experiments with emojis a symbolic system, **we used the name of the emoji and its representation together as input for the vector embedding** (columns `Name` and `Emoji` from the dataset). Initially we wanted to use only the representation, but this did not produce useful results with our choice of embedding model. Also, it must be noted that we had to bundle together the emojis with different skin tones into one entity.

Using the aforementioned datasets, obtaining the embeddings and making a histogram of what motifs from TMI are associated to each card, we obtain Fig. 6, which could either be telling that folk tales are overwhelmingly close to the semantics of the Ace of Swords, or we might be looking at some strange bias that needs too be corrected. In any case, besides these being interesting research questions on their own that we will get to in the full publication, since the TMI has over 46 thousand items this is not such a problem, as we can

always pick *balanced* subsets (i.e. a fairer distribution of card bins) of hundreds of motifs, which is enough for most game design needs. We have developed a few simple algorithms that try to optimize balancing between multiple symbolic systems, even, in the case this is a requirement for a specific scenario (i.e. have a set of motifs that is both evenly distributed among all Tarot cards *and* hexagrams).

In our experiments many things affect the final binning (of motifs into cards, hexagrams) and relative cosine distances, most notably the choice of text used as the input for the embedding (e.g. variations of the full json entries in the datasets), but it is still an open empirical question if this is a problem for doing game design in the long run. Our running hypothesis is that it is not, but to be able to test it we need to develop theoretical and technical tools that allow us to better analyze what is happening within these datasets and their latent spaces. We encourage the reader to try to reproduce our experiments in any way they can and see what they can get out of it. Why not just throw the text of this paper into your favorite *vibe coding* environment and see if it can reproduce this basic setup?

B A SHORT AND GENTLE INTRODUCTION CATEGORY THEORY FOR PEOPLE WHO KNOW (A BIT) ABOUT GRAPHS

Category theory (CT) is a **famously arcane** and indigestible field of mathematics, considered by many “too abstract”. From my own personal journey with CT, part of the issue is that it oscillates quickly between trivially simple (its just arrows!) to extremely complex and formal, with many terms that are *somewhat* precise but also allegoric, such as *sheaves*, *stalks*, *fibers*, *pullbacks*, *cones*, *horns*, etc. At the same time, most material on the subject stays at an abstract level and provide little examples or instantiations of such objects that might help the reader bring them to something concrete. Rosiak’s book “Sheaf Theory Through Examples” [37] was a major breakthrough (get the full version if you can) in bridging the theoretical with the practical, and it would be my go-to starting literature for anyone interested.

In this section I wish to give some gentle, bare-minimum introduction that might be useful for the reader to understand our main text, and some interesting philosophical implications of CT. For a computer scientist, it is useful to **import from** all their knowledge of graphs and networks, so you can feel more at ease in this “turf”. Although not all graphs are categories, all categories have an underlying (directed) graph that can be extracted, and many of the methods of working within CT are graph-based with an extra degree of mathematician neuroticism. Almost all introductions to CT start giving up the same bare-bones definition, which should be our starting point. *A category is a:*

- **Collection of objects**, which in our graph perspective are vertices, but they can be *things* in the most abstract sense, including other categories;
- **A set of arrows**, called *morphisms*, which are essentially directed edges and indicate there is a relationship between these two objects. For instance, in the category $\mathbb{N}$ with morphisms given by the squaring operation, there is an arrow $3 \rightarrow 9$;
- **A rule for composing arrows**, so that if there is an arrow (morphism) $f: A \rightarrow B$ and an arrow $g: B \rightarrow C$, then there is a composite morphism $g \circ f: A \rightarrow C$. This is just like saying valid paths in the graph can be treated as just one arrow; and
- **An (unique) identity morphism for each object** (id_A , id_B , etc.), with is just an arrow that has the same object as its source and target, a loop, a *do-nothing* of sorts.

And all of these must also follow two other rules:

- **Associativity**, which means you can compose in any order, e.g. $h \circ (g \circ f) = (h \circ g) \circ f$; and




Figure 5: Simple illustration of category from a graph: A, B, C and D are our objects or vertices, and $f: A \rightarrow B, g: B \rightarrow C, h: C \rightarrow D, i: D \rightarrow B$ our morphisms or edges. This gives us the graph of the left with only the black arrows. Then, we define that morphisms can be composed by concatenating paths (pink and dark yellow arrows). Finally, we also have one identity morphism for each object in the form id_X (green arrows). An interesting implication of CT is that, because the identity morphisms are unique to each object, they can actually stand *for them*, and we can completely ditch the object set and work only with morphisms (right), and still be able to produce the same results.




Figure 6: Binning of the whole TMI dataset within the 56 minor arcana cards (i.e. what is the card closest to each motif). Notice how the overwhelming majority of cards are closer to the King of Swords. This could just be an issue with the pipeline used, or indicate some deeper bias in the data.

- **Unit laws**, which assert that for any morphism $f: A \rightarrow B$, one has $id_B \circ f = f$ and $f \circ id_A = f$. In other words, when the identity morphism is part of a path, it does not do anything (which is intuitive since its just a self-loop).

If you have all of those things, you have a category! Fig. 5 has a nice graphical demonstration of all of these put together for a simple category with four objects and four morphisms. Now the reader might be asking *how can these simple rules guarantee the super powers attributed to Category Theory?* And the answer to that is that both the objects and the arrows can be basically anything, including arrows and other categories. And, if you follow these rules (composition, associativity, having an identity, etc.) you guarantee the preservation of **structure** (in an algebraic and in most cases topological sense), so that you can build complexity from simplicity by stacking and nesting objects and arrows ad infinitum.

The whole field of CT is basically mathematicians getting carried away with this and finding different ways these constructions hold across different parts of mathematics. And while most of mathematics *can* be organized into categories, any structure that breaks one of the axioms can't play along. For example, octonions—an eight-dimensional extension of the quaternions—aren't associative ($A(BC) \neq (AB)C$), so they refuse to form a strict category (but quaternions are quite happy to!). CT has a rich (if not overwhelming)




Figure 7: Sampling of 100 motif beads balanced by the 56 bins (cards) of the Tarot. Compare to the unbalanced sampling of Fig. 1, where swords (left) are overwhelmingly represented, which is a consequence of the global distribution of cards seen in the histogram of Fig. 6.

dialect within mathematics full of idiosyncrasies and visual thinking tools, such as the [snake lemma](#), and [diagram chasing](#), which is an interesting method for finding proofs diagrammatically [5]. This is also one of the reasons we believe vis should have a closer rapport with CT.

In our proposed approach, we can start to frame it within CT by considering every "set" (motifs, Tarot cards, I Ching) as a discrete category, which is a category with just objects and their identity morphisms (just elements and no extra behavior). Then, we map these categories to "enriched" versions with the chosen vector embeddings. In CT, we call operators that map categories into other

categories **functors**, and they are a very powerful concept because they are able to reassign all objects and arrows from one category to another while preserving structure. In this case this is trivial because our discrete categories just have the identity morphisms. Once we have the embedded categories, we can have morphisms from the motifs into the symbolic systems as **posets** (ordered sets), which is what the Fig. 3 center and right parts are showing. From there on, there are many methods that one can use to compare motifs based on their orderings within a symbolic system, and between different symbolic systems.

A final interesting thing to point out about CT is that it has an essentially *Buddhist* metaphysics. While it might seem from our usual way of seeing things that the *objects* are what matter in the category, this could not be farther from the truth. Because we are required to have identity morphisms for each object, all information we need is contained within the set of morphisms and we can essentially discard the objects (Fig. 5, right). It is as if all things are *void of essence*, as if they are actually defined only by their relationships with other things. This idea is *isomorphic* to the Buddhist concept of *Sunyata*. A stronger version of this perspective is expressed within CT by the *Yoneda Lemma*, which we previously mentioned. It is a common misconception that when *Buddhism* talks about “emptiness”, it is denying existence on a prescriptive sense, as if “nothing exists” or reality is empty. Rather, it is about the impossibility of talking about the essence of “things” that are, by all means, defined by their relationships rather than by some fundamental essence. This supposed “essence” is also isomorphic to Kant’s inaccessible *noumena*, which certainly inspired Joseph Knecht’s epiphany we quoted at the end of the paper.