TRANSFORM: Advancing ideational social science through transformer-based (language) models: a pilot study on the imagined futures of Web3.

Timo Seidl (University of Vienna)

'Ideas are primary exhibits for the doctrine that what is important to study cannot be measured and that what can be measured is not important to study.'

— Philip Converse

I. The Study of Ideas in the Social Sciences

Ever since the days of Max Weber, social scientists have grappled with the question of **how ideas can become 'effective forces in history'** (Weber 1920/2007, 48). However, despite their illustrious social scientific pedigree, ideas have **long lived a life in the shadows of interests and institutions**.

Two factors inhibited the development of ideational explanations.

- First, ideas were **under-conceptualized** and their causal status thus remain **under-specified**.
- Second, ideas were **under-operationalized**.

However, over the last quarter of a century, this has changed (see Figure 1). The 'ideas school' (Kamkhaji and Radaelli 2022) in the social sciences has made great inroads into more carefully conceptualizing ideas and clarifying their causal role in social and political processes (e.g., Berman 2013; Blyth 2003; Schmidt 2008). These **conceptual advances** have given rise to a flurry of empirical studies which have, however, been mostly qualitative.

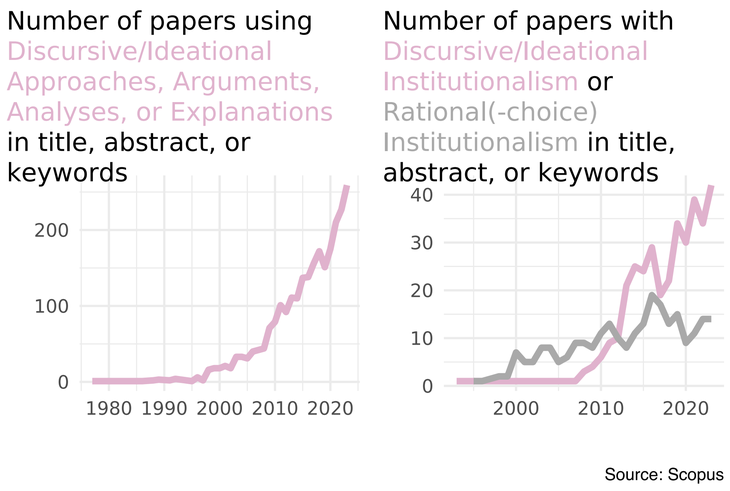

Figure 1

What is often missing from this literature are more systematic attempts to **measure ideas**. In other words, what is missing are quantifiable indicators of text that allows scholars to **operationalize ideational concepts** such as narratives or frames **in a scalable manner**. In other words, what is missing are approaches where the marginal human input declines over time as machines learn to reliably replicate human tasks such as classification.

II. The Computational Turn in Text Analysis

In parallel with the formation of the ideas school, the growing availability of digital (text) data and the rapid improvements in computational power have given rise to a methodological movement that treats **text as data** (see Figure 2).[1] Often adapting approaches from computer sciences or linguistics for social scientific purposes, this literature has developed a suite of 'tools, models, and software that **facilitate[s] the analysis and organization of texts at scale**' (Grimmer, Roberts, and Stewart 2022, 3; see also Ash and Hansen 2023; Gentzkow, Kelly, and Taddy 2019; Van Loon 2022).

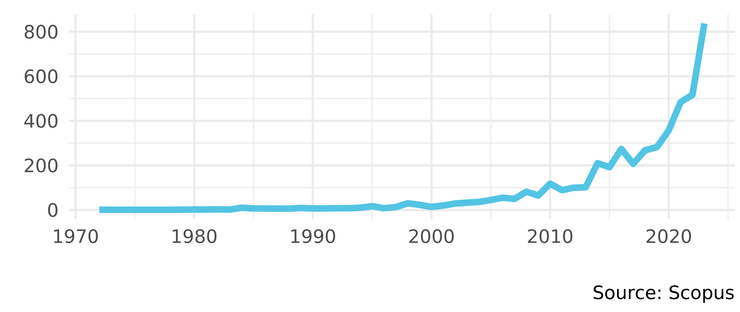

Figure 2

Much of the early text-as-data literature relied on **bag-of-words approaches**, i.e., for each document they count how many times each word appears in it. While undoubtedly powerful, these approaches also throw out a lot of information—most notably word order—and generally result in very sparse document-feature matrices (since most documents do not contain most words).

More recently, scholars have started to use **distributed representations** which represent words as dense vectors in a relatively low-dimensional vector space and thus encode similarity (Rodriguez and Spirling 2022). Crucially, this is a form of **transfer learning** as these vectors (or 'embeddings') are learned from large amounts of unlabeled data (Grimmer, Roberts, and Stewart 2022, 79).

The same is true for transformer-based large language models, which provide **contextual word embeddings** that not only capture the meaning of words but also their relation to other words. Their central advantage is also that they leverage their **general language and task-specific knowledge** for context-specific problems such as text classification. This generally allows them to do more complex tasks or do existing tasks better (Laurer, Van Atteveldt, et al. 2023; Widmann and Wich 2022). In addition, they can handle very different input data due to being **multi-modal** and **multilingual** (e.g., Licht and Lind 2023).

III. Imagined Futures & Web3

So far, we have documented both **conceptual and technical advances** in how to understand ideas and measure text. It is the **core proposition of this project** that there is great potential in bringing these advances in ideational and computational social science together. This is a big claim, and the proof of the pudding is in eating. Therefore, the **core idea behind this project** is to lend plausibility to this promise through a pilot study on imagined futures and Web3.

There are **four central assumptions** behind the ideational concept of imagined futures (Beckert 2016; Beckert and Bronk 2018; cf. Jasanoff and Kim 2015):

> **i** The Four Assumptions of the Imagined Futures Approach
>
> 1. Whatever their goals are, actors need to act based on expectations of presumed outcomes. This means that how they imagine the future shapes their present actions.
> 2. The future is often radically uncertain and thus genuinely unknowable. This means that actors cannot rely on rational expectations but instead need to resort to 'fictional expectations'.
> 3. Fictional expectations about future states of the world are not formed in isolation but are socially constructed. This also means that they can be reconstructed through discursive trace data.
> 4. Fictional expectations are not only intersubjectively developed, they can also affect social processes. This means that reconstructing them can help us understand real-world dynamics in politics, economics, and culture.

Web3—the idea to rebuild the web around public blockchains and token-based economics—is a **particularly promising case** to study the social construction and consequences of fictional expectations. As both journalists (McKenzie and Silverman 2023) and academics (Shiller 2019) have already pointed out, the periodic rise and fall of Web3 and cryptocurrencies in particular are **intimately tied to storytelling and the personal and collective futures its proponents and critics have advanced**. There is also an enormous amount of discursive trace data, from newspapers and blog posts to memes and videos, that makes Web3 a particularly apt case study.

IV. Empirical Approach

This project wagers that **transformer-based language models hold the promise to help us study complex ideational concepts at scale**. They can do so by aiding researchers in the three '**core tasks of social science research**' (Grimmer, Roberts, and Stewart 2022) summarized in Figure 3.

Figure 3

i. Discovery

Discovery is about refining or creating new conceptualizations or ways of organizing the world (Grimmer, Roberts, and Stewart 2022). Figure 4 illustrates how this process of concept development and refinement can be thought of as iterating between theoretical, substantive, computational and interpretative work (see also Nelson 2020; Carlsen and Ralund 2022).

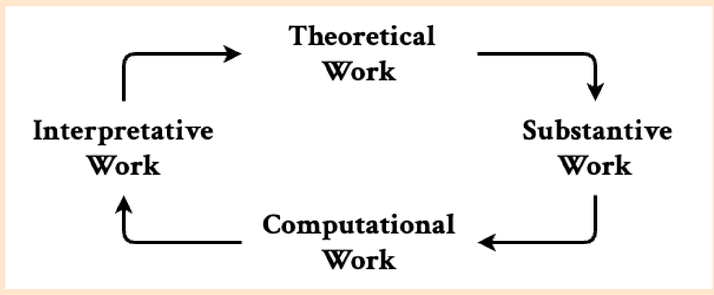

Figure 4

In the context of this project, discovery will be relatively **deductive** as we start from the theoretical concept of fictional expectations. However, while this concept is fairly clearly defined, it will still be necessary to identify **different and competing fictional expectations about the future of Web3** (e.g. personal enrichment, decentralizing finance, building a new type of society) (cf. Gehring, Adema, and Poutvaara 2022). Doing so requires close engagement with the substantive literature on Web3. But it can also be aided unsupervised computational methods such as BERTopic (Grootendorst 2022). Such transformer-based methods provide insightful summaries of large amounts of text and can be used to amend the list of fictional expectations identified by the literature.

ii. Measurement

Measurement is about connecting concepts to data, allowing scholars to understand the prevalence and evolution of these concepts in the real world. Much of the excitement around computational text analysis stems from the fact that 'it greatly expands the range of phenomena that can be tractably quantified' (Grimmer, Roberts, and Stewart 2022, 171). A key way in which transformer-based models can improve on existing methods is when it comes to complex classification tasks. This can help **bridge the gap between sophisticated theoretical concepts and the empirical proxies used to measure these concepts** (Baden et al. 2022).

In this project, a central goal is to classify text into different substantive categories of fictional expectations. A first task, given that fictional expectations 'delineate visions (of the future) that go beyond observable truths' (Beckert and Bronk 2018, 10), is to **identify texts that are about the future** (see also Müller 2022; Müller and Proksch 2023). A second step is to identify those future-oriented texts that are about Web3. Lastly, based on the categories developed during discovery, we can further group these fictional expectations about the future of Web3. This process, **illustrative of several potential strategies** that will be pursued, is summarized in Figure 5.

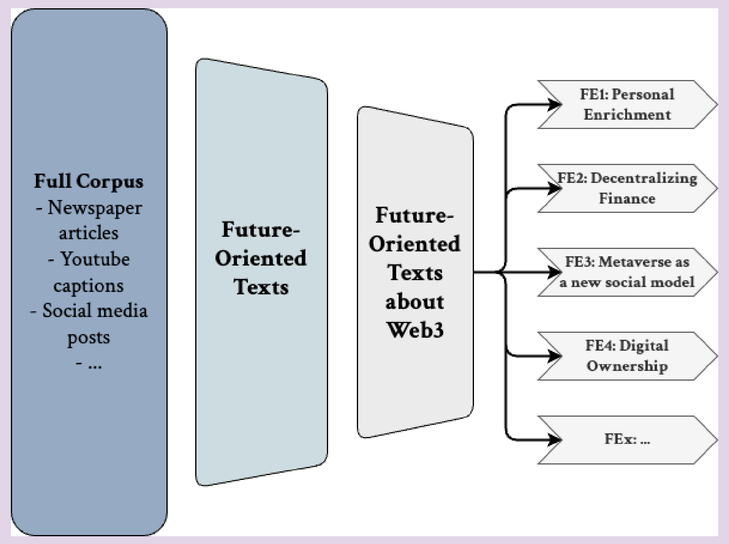

Figure 5

All of these tasks are separate classification tasks. Each of them therefore requires **careful validation through substantive, structural, and external evidence**—following the suggestions by the *ValiTex* framework (Birkenmaier, Wagner, and Lechner 2023; cf. Baden et al. 2022). This is not just important to ensure that the measures actually measure what they portend to measure, but also to ensure that the models do not harbor **biases** related to, for example, gender or race. Different classification strategies and models will be compared, ranging from large proprietary generative models able to perform zero-shot or few-shot classification to smaller, open-source models fine-tuned for a specific task (see also Laurer, van Atteveldt, et al. 2023; Törnberg 2024).

iii. Inference

Inference is about making predictions or causal inferences about human behavior based on these measures. While measurement is crucial, the goal should be not just a descriptive mapping of different fictional expectations, but also to use them to explain real-world political, economic or social processes. Figure 6 explores, in a very stylized fashion, two potential avenues for using fictional expectations to predict, respectively, the valuations of cryptocurrencies or Web3 start-ups or regulatory outcomes.

For example, this could involve running a time-series regression with fictional expectations being the independent variable and currency valuation being the dependent variable, with 'hard' economic variables as a control. *In other words, if 'valuation is expectation and expectation is imagination' (Shackle 1972/1992, 8), does sentiment and/or category of fictional expectations predict valuations independent of different economic control variables?*. Another example would be to look at how fictional expectations spread. *Do legislators adopt certain fictional expectations from private actors and does this predict regulatory outcomes?*

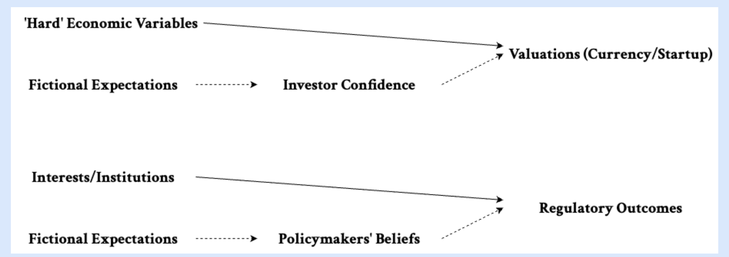

Figure 6

> Tip
>
> *For a pdf version including a complete list of references, scan the above QR code!*

[1] Text here, while indeed central to this project, is a *pars pro toto* for a wider set of symbolic information including images, sound, or video.

V. References

Ash, Elliott, and Stephen Hansen. 2023. "Text Algorithms in Economics." *Annual Review of Economics* 15 (1): 659–88. https://doi.org/10.1146/annurev-economics-082222-074352.

Baden, Christian, Christian Pipal, Martijn Schoonvelde, and Mariken A. C. G van der Velden. 2022. "Three Gaps in Computational Text Analysis Methods for Social Sciences: A Research Agenda." *Communication Methods and Measures* 16 (1): 1–18. https://doi.org/10.1080/19312458.2021.2015574.

Beckert, Jens. 2016. *Imagined Futures: Fictional Expectations and Capitalist Dynamics*. Cambridge, Mass.: Harvard University Press.

Beckert, Jens, and Richard Bronk, eds. 2018. *Uncertain Futures: Imaginaries, Narratives, and Calculation in the Economy*. Oxford: Oxford University Press.

Berman, Sheri. 2013. "Ideational Theorizing in the Social Sciences Since 'Policy Paradigms, Social Learning, and the State'." *Governance* 26 (2): 217–37. https://doi.org/10.1111/gove.12008.

Birkenmaier, Lukas, Claudia Wagner, and Clemens Lechner. 2023. "Vali-Tex – a Unified Validation Framework for Computational Text-Based Measures of Social Science Constructs." https://arxiv.org/abs/2307.02863.

Blyth, Mark. 2003. "Structures Do Not Come with an Instruction Sheet: Interests, Ideas, and Progress in Political Science." *Perspectives on Politics* 1 (4): 695–706. https://doi.org/10.1017/S1537592703000471.

Carlsen, Hjalmar Bang, and Snorre Ralund. 2022. "Computational Grounded Theory Revisited: From Computer-Led to Computer-Assisted Text Analysis." *Big Data & Society* 9 (1): 205395172210801. https://doi.org/10.1177/20539517221080146.

Gehring, Kai, Joop Adema, and Panu Poutvaara. 2022. "Immigrant Narratives." https://doi.org/https://www.cesifo.org/DocDL/cesifo1_wp10026.pdf.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy. 2019. "Text as Data." *Journal of Economic Literature* 57 (3): 535–74. https://doi.org/10.1257/jel.20181020.

Grimmer, Justin, Margaret E. Roberts, and Brandon M. Stewart. 2022. *Text as Data: A New Framework for Machine Learning and the Social Sciences*. Princeton: Princeton University Press.

Grootendorst, Maarten. 2022. "BERTopic: Neural Topic Modeling with a Class-Based TF-IDF Procedure." arXiv. https://arxiv.org/abs/2203.05794.

Jasanoff, Sheila, and Sang-Hyun Kim, eds. 2015. *Dreamscapes of Modernity: Sociotechnical Imaginaries and the Fabrication of Power*. Chicago: The University of Chicago Press.

Kamkhaji, Jonathan C., and Claudio M. Radaelli. 2022. "Don't Think It's a Good Idea! Four Building Sites of the 'Ideas School'." *West European Politics* 45 (4): 841–62. https://doi.org/10.1080/01402382.2021.1959751.

Laurer, Moritz, Wouter Van Atteveldt, Andreu Casas, and Kasper Welbers. 2023. "Less Annotating, More Classifying: Addressing the Data Scarcity Issue of Supervised Machine Learning with Deep Transfer Learning and BERT-NLI." *Political Analysis*, June, 1–17. https://doi.org/10.1017/pan.2023.20.

Laurer, Moritz, Wouter van Atteveldt, Andreu Casas, and Kasper Welbers. 2023. "Building Efficient Universal Classifiers with Natural Language Inference." arXiv. https://arxiv.org/abs/2312.17543.

Licht, Hauke, and Fabienne Lind. 2023. "Going Cross-Lingual: A Guide to Multilingual Text Analysis." *Computational Communication Research* 5 (2): 1. https://doi.org/10.5117/CCR2023.2.2.LICH.

McKenzie, Ben, and Jacob Silverman. 2023. *Easy Money: Cryptocurrency, Casino Capitalism, and the Golden Age of Fraud*. New York: Abrams Press.

Müller, Stefan. 2022. "The Temporal Focus of Campaign Communication." *The Journal of Politics* 84 (1): 585–90. https://doi.org/10.1086/715165.

Müller, Stefan, and Sven-Oliver Proksch. 2023. "Nostalgia in European Party Politics: A Text-Based Measurement Approach." *British Journal of Political Science*, November, 1–13. https://doi.org/10.1017/S0007123423000571.

Nelson, Laura K. 2020. "Computational Grounded Theory: A Methodological Framework." *Sociological Methods & Research* 49 (1): 3–42. https://doi.org/10.1177/0049124117729703.

Rodriguez, Pedro L., and Arthur Spirling. 2022. "Word Embeddings: What Works, What Doesn't, and How to Tell the Difference for Applied Research." *The Journal of Politics* 84 (1): 101–15. https://doi.org/10.1086/715162.

Schmidt, Vivien A. 2008. "Discursive Institutionalism: The Explanatory Power of Ideas and Discourse." *Annual Review of Political Science* 11 (1): 303–26. https://doi.org/10.1146/annurev.polisci.11.060606.135342.

Shackle, G. L. S. 1972/1992. *Epistemics & Economics: A Critique of Economic Doctrines*. New Brunswick, N.J.: Transaction Publishers.

Shiller, Robert James. 2019. *Narrative Economics: How Stories Go Viral & Drive Major Economic Events*. Book Collections on Project MUSE. Princeton: Princeton University press.

Törnberg, Petter. 2024. "Best Practices for Text Annotation with Large Language Models." arXiv. https://arxiv.org/abs/2402.05129.

Van Loon, Austin. 2022. "Three Families of Automated Text Analysis." *Social Science Research* 108 (November): 102798. https://doi.org/10.1016/j.ssresearch.2022.102798.

Weber, Max. 1920/2007. *The Protestant Ethic and the Spirit of Capitalism: Translated by Talcott Parsons. With an Introduction by Anthony Giddens*. London: Routledge.

Widmann, Tobias, and Maximilian Wich. 2022. "Creating and Comparing Dictionary, Word Embedding, and Transformer-Based Models to Measure Discrete Emotions in German Political Text." *Political Analysis*, June, 1–16. https://doi.org/10.1017/pan.2022.15.