

AUGUST 2020

Mainframes: A Provisional Analysis of Rhetorical Frames in AI

CSET Issue Brief



AUTHORS

Andrew Imbrie
James Dunham
Rebecca Gelles
Catherine Aiken

Executive Summary

When it comes to artificial intelligence, the headlines suggest that great powers are engaged in an AI arms race: “For Superpowers, Artificial Intelligence Fuels New Global Arms Race,” reads one story in *Wired*.¹ “China is Winning a New Global Arms Race,” observes *Bloomberg Markets and Finance*.² One report in *The Wall Street Journal* asserts, “The New Arms Race in AI.”³ To what degree do these headlines accurately represent elite opinion about AI?

Framing technological competition as an “AI arms race” or “battle for supremacy” has implications for policy, security, and international cooperation.⁴ Maintaining U.S. leadership in strategic areas of artificial intelligence is an important policy objective, but AI is a general-purpose technology that enables a vast array of applications. The AI arms race narrative obscures more than it reveals about which applications deserve priority and how the United States can shape norms and standards to guide the development of AI and machine learning consistent with democratic principles.⁵ Perceptions of an arms race in AI could lead great powers to disregard investments in safety and security. Further, arms race dynamics may inhibit cooperation and increase the risk of miscalculation or misperception involving the use of AI-enabled platforms and capabilities.⁶

Are the great powers engaged in an AI arms race? Emerging patterns in mass communication provide insight. Opinion makers use language to influence the behavior of others and justify new investments and directions in policy.⁷ Public perceptions of AI shape the calculus of national leaders. Rhetorical framing of AI therefore serves as a barometer for both public and elite opinion, indicating whether attitudes are becoming more cooperative or competitive over time.

In this paper, we develop a novel methodology to explore the rhetorical framing of AI. We searched more than 4,000 English-language articles over the 2012 to 2019 period from the global news outlet Reuters, the U.S.-based *Defense One*, and *Foreign Affairs* and identified references to “AI competition.”⁸ This frame describes AI development as a race between two or

more actors, such as governments or companies—whether in military or non-military terms. Invocations of the frame include the following:

- a military competition (“arms race”)
- historical competition (“Cold War” or “Sputnik Moment”)
- a territorial competition (“supremacy in Europe”)
- a competition for resources (“battle for talent”)
- any other type of competition (“two-man contest” or “AI rivalry”)

In addition to capturing occurrences of the frame over time, we identified the purpose for which it is invoked. We classified the purpose of the “competition” frame as *motivation* (for changing the status quo), *explanation* (of the state of the world), or *critique* (of the frame itself).

Our provisional findings are as follows:

- **Since 2012, a growing number of articles in the three news sources have included the competition frame, but prevalence of the frame as a proportion of all AI articles peaked in 2015.** Reporting on AI has become more sophisticated and diverse.
- **The competition frame was predominantly used for explanation, though an increasing share of articles invoked it for motivation or critique.** The purpose seems closely related to the source: motivation and critique were more common in *Defense One*, which features analysis and opinion pieces together with news, than in Reuters, a news agency.
- **In the United States, individuals most frequently associated with use of the frame for motivational purposes came from the tech sector, the U.S. Senate, and the U.S. Department of Defense.**
- **Internationally, the institutional affiliations most frequently connected with use of the frame for motivational purposes were the French and Russian governments.**
- **The companies most frequently associated with the frame were from the United States, Japan, Finland, India, and China.**

Introduction

World leaders are framing advances in artificial intelligence among major powers as an AI arms race. In 2017, Russian President Vladimir Putin remarked to a group of students, “Whoever becomes the leader in this sphere will become the ruler of the world.”⁹ One year later, Chinese President Xi

Jinping declared, “Accelerating the development of a new generation of AI is an important strategic handhold for China to gain the initiative in global science and technology competition.”¹⁰ In June 2019, the EU High-Level Expert Group on Artificial Intelligence issued recommendations on AI policy and investment, urging greater prioritization and swift action: “With the digital economy being characterised [sic] by a ‘winner takes all’-logic and diminishing returns for late entrants, it is essential for European actors to invest as first movers so as to have the possibility to capture large market shares.”¹¹ The Trump administration is pursuing an equally ambitious agenda on AI. According to its Chief Technology Officer, “Our goal is very clear, the unique American ecosystem must do everything in its collective power to keep America’s lead in the AI race and build on our success.”¹²

To what extent do these statements reflect elite opinion on AI? Different sources may reveal different patterns about the direction of the field. Remarks or speeches from national leaders will depart in their meaning and tone from news wires, online commentary, social media posts, and blogs. We include in this analysis three news sources – Reuters, *Defense One*, and *Foreign Affairs* – that vary in their content and audience. We searched these sources for articles published between 2012 and 2019 that mention “artificial intelligence” or “AI,” which yielded 4,009 articles.¹³

We focus on the framing of AI development as a competition, either between nations or private companies. The canonical invocation of this rhetorical frame is “the AI arms race,” while other military references describe AI development as a cold war battle. We also consider allusions to a “race,” “rivalry,” or “contest” to be examples of the competition frame.

In addition to the incidence of the rhetorical frame, we identified three purposes for using the frame:¹⁴

- **Explanation:** to characterize the state of the world and describe events, including events that may occur in the future, and events that are dangerous or could lead to instability and conflict, such as arms races.
- **Motivation:** to justify new policies, investments, or force postures that will allow one side to maintain a competitive edge or to “win” the race. In some cases, the frame is invoked to justify new policies and investments that will have positive-sum effects and benefit all sides.
- **Critique:** to condemn the use of the frame itself. When the purpose is critique, there is an explicit rejection of the frame, not just a rejection of the world or reality it describes. This option includes language

about the frame and why using it can fuel negative dynamics, instability, escalation risks, or result in unintended consequences.

We also identified examples of reporting on structured competition, like the Google DeepMind Challenge Match, but did not consider these invocations of the competition frame, so they were excluded from further analysis.

Data and Measurement

To explore the use of the AI competition frame from 2012 to 2019, we analyzed AI articles from three different news sources that varied in content, coverage, and audience. We included articles from Reuters news agency for broad, global news coverage; *Defense One* for more specialized opinion and analysis about U.S. national security; and *Foreign Affairs* for longer-form coverage of international relations. We selected these sources based on two criteria: first, they offered a diversity of news analysis, commentary, and opinion; and second, the texts were more consistently structured (and therefore amenable to machine analysis) than speeches, blogs, or social media posts. In future extensions, we intend to investigate whether our findings hold at scale and for other mediums.

Reuters is the most prolific source and returned 3,436 AI articles, while *Defense One* and the bi-monthly *Foreign Affairs* returned 518 and 55, respectively.¹⁵ Coverage of AI in all three sources greatly increased between 2012 and 2019. For example, AI was mentioned just 26 times across Reuters articles in 2012 compared to 1,106 times in 2019. The first mentions of AI in *Foreign Affairs* and *Defense One* were in 2013 and 2014, respectively.¹⁶ In 2019, *Defense One* published 215 such articles and *Foreign Affairs* 19.

A prerequisite for analysis of this rhetorical frame is the ability to measure its use accurately and consistently in all contexts. To identify use of the competition frame in our selected sources, we relied on manual annotation, following a carefully developed annotation framework (Appendix B). At least one annotator coded each article for whether it included the frame, and if so, its apparent purpose: explanation, motivation, critique, or structured competition. Two annotators coded 10 percent of articles, independently, to enable tracking of inter-coder agreement. They also extracted any individuals or institutions mentioned in connection with the frame, typically through quotation in the Reuters articles.

To assess how well we defined the concepts of interest and applied these definitions in annotation, two annotators coded 10 percent of the articles. Each annotator was unaware of the existence of the other annotator's work to ensure the annotations remained independent. These double annotations were then assessed for inter-coder agreement. For document-level annotations, the two annotators agreed on all tags 98 percent of the time. For paragraph-level annotations, annotators agreed on all tags 71 percent of the time, with 50 percent of the disagreements coming from formatting variations like capitalization, punctuation, and truncation rather than disagreements on specific codes. Given the high level of inter-coder agreement found in the 10 percent of articles annotated twice, we determined that annotators were capable of coding independently.

Ascertaining the frequency of use and purpose of the frame is challenging for several reasons:

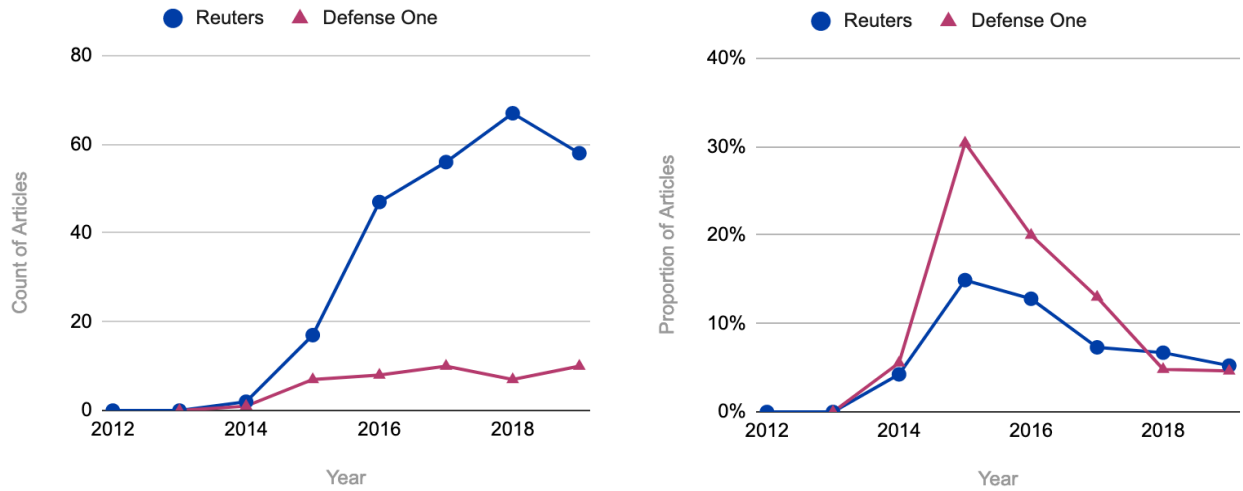
- **Tractability:** rhetorical frames are fairly abstract concepts, and defining one precisely enough to identify its invocation can be difficult. To overcome this challenge, we developed clear definitions, annotation guidelines, and examples to clarify exactly what was meant by each frame purpose.
- **Consistency:** the competition frame needs to be identified in a consistent way over time and across observers. To mitigate this challenge, we provided annotators with clear guidelines. We also tracked inter-coder agreement to ensure the guidelines were applied consistently between annotators.
- **Inferring intent:** our analysis of the competition frame's purpose requires that we divine speakers' implicit intentions, which may be more or less apparent. This challenge was also mitigated by using multiple annotators and tracking inter-coder agreement. If multiple annotators infer the same intent, it is likely the correct perceived intent is being inferred.
- **Potential confounders:** artifacts of collection or classification should not affect our inferences about patterns in frame use or purpose. For this study, we consider only three news sources. In future work, we intend to consider a wider variety of sources.

Analysis

Since 2012, a growing number of articles have included the competition frame, but prevalence of the frame as a proportion of all articles about AI

peaked in 2015 (Figure 1). We expect that as the applications of AI expanded during this period, reporting on AI became more diverse. As there was only one use of the competitive frame in the *Foreign Affairs* dataset, we considered only Reuters and *Defense One* below.

Figure 1. Number and proportion of AI articles with competition frame over time.¹⁷



Reuters articles are often straight news, while *Defense One* provides more opinion and commentary. We see this pattern in the frequency of each purpose for the competition frame. In Reuters, explanation predominates as the frame purpose. Of 389 invocations of the competition frame in Reuters articles, 352 (90 percent) are for explanation. The motivation and critique frames are more clearly represented in the *Defense One* data. Of 79 instances of the frame in *Defense One* articles, 46 (58 percent) are for explanation. There is little variation in frame purpose in Reuters between 2012 and 2019 (Figure 2). Meanwhile, in *Defense One*, critique becomes more prominent as a frame purpose in 2017, while motivation reaches peak use in 2018 (Figure 3).

Figure 2. Frame purpose in Reuters over time.





Figure 3. Frame purpose in Defense One over time.



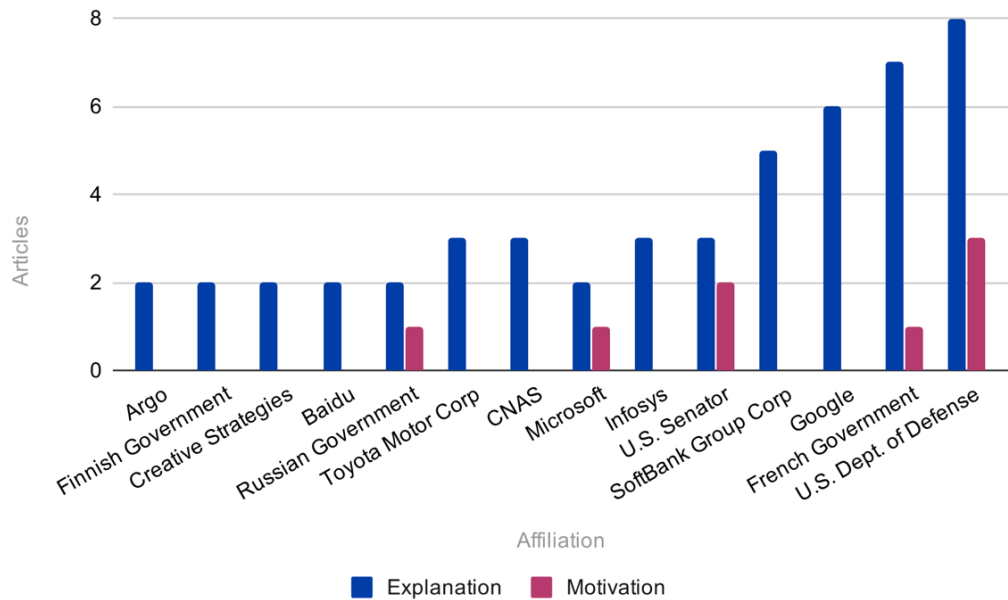
We also analyzed the origins and prominence of the individuals who are most frequently associated with use of the competition frame.¹⁸ Those most often associated with employing the frame for explanatory purposes were a former senior U.S. defense official, a European head of state, an auto executive in Japan, a Japanese-Korean technology entrepreneur, and a technology columnist. That a range of figures appear during invocations of the frame for explanatory purposes suggests the competition frame has diffused widely in public discourse and become shorthand for understanding the larger geopolitical context of investments in AI.

In the United States, individuals most often quoted during uses of the frame for motivational purposes came from the U.S. tech sector, the U.S. Senate, and the U.S. Department of Defense. Although provisional, these findings suggest that national security institutions are involved in employing the frame to mobilize latent resources, increase budgets, and drive breakthroughs in strategically relevant technologies.

The companies most frequently associated with the frame were from the United States, Japan, India, and China. The United States accounts for 65 percent of the world’s top AI start-ups, with China coming in second at 6 percent.¹⁹ That companies in Japan and India are associated with the frame suggests that venture capital AI investments and deal counts will continue to grow more diverse.

Internationally, affiliation data is sparse for those who were most frequently associated with use of the frame for motivational purposes. Of those that we identified, the two most common affiliations were the French and Russian governments. Analysts have suggested that leaders of both governments see AI through a geopolitical lens and competition will likely ramp up if present trends hold.²⁰ Considering data limitations, we leave assessment of this possibility to future research.

Figure 4. Frame purposes of articles with most common affiliations of individuals associated with the frame.²¹



The proportion of articles with the competition frame increased after 2012 and peaked in 2015. These dynamics are consistent with broader trends in the field. In 2010, the computer scientist Fei-Fei Li launched the ImageNet Large Scale Visual Recognition Challenge, a database of 14 million labeled images that enabled developers to improve the accuracy and performance of machine learning algorithms. In the 2012 ImageNet Challenge, the neural network AlexNet achieved a remarkable feat: an error rate of 15.3 percent, besting the nearest competitor by more than 10 percentage points.²² Supervised learning gained prominence in the field, and the extraordinary pace of progress drew the interest of major powers, including the United States and China. The peak in the competition frame in 2015 may reflect intensifying geopolitical dynamics; its subsequent decline may indicate that assessments of the field are growing more diverse and sophisticated. Future extensions of this project may qualify or alter these findings. Reuters is a newswire that features analysis and reportage. Less structured texts, policy outlets, and sources with more opinion writing could reveal different trends.

In addition to tracking the frequency of use, this analysis sheds light on the purpose for which the competition frame is invoked. Explanation predominates, although a rising share of articles invoke the frame for motivation or critique. Some actors invoke the frame to mobilize latent resources and increase budgets, while others react critically to the emerging competitive discourse.

CSET plans to continue its analysis of the competition frame across different venues for public discourse and corpora that are less structured and more opinion-driven, including governmental and international reports, floor statements and speeches, popular press such as *Wired*, blogs, commentary on news wires, and social media. Building on our analysis, future research could evaluate additional variables, such as origins (who initiated use of the frame in a specific context or sector), prominence (citation counts and network analysis to determine impact), and mode-switching (invoking the frame for different purposes to different audiences). Additional research is needed on the rhetorical frames employed in national AI strategies. It would also be instructive to compare the competition frame with other rhetorical frames, such as those depicting AI as a “flash in the pan” or “economic gold rush,” as well as those warning of “killer robots” and “a world without work.” Future extensions could analyze the use and evolution of these frames in other countries limited to the last two decades of digitized data.

Acknowledgments

The authors are grateful to Philippe Loustaunau and Linnea Micciulla for their contributions to the data annotation process. Michael C. Horowitz, Director of Perry World House and Richard Perry Professor at the University of Pennsylvania, and Michael Sellitto, Deputy Director of the Stanford Institute for Human-Centered Artificial Intelligence, provided excellent feedback. Many thanks to Ben Buchanan, Dewey Murdick, Igor Mikolic-Torreira, Alexandra Vreeman, and Lynne Weil for their comments and editorial support. Kady Arthur provided valuable research assistance. All errors are the authors' own.



© 2020 Center for Security and Emerging Technology. This work is licensed under a Creative Commons Attribution-Non Commercial 4.0 International License.

To view a copy of this license, visit:
<https://creativecommons.org/licenses/by-nc/4.0/>.

Appendix A: Supplementary Table

Table A.1: Counts of articles analyzed for AI and competitive frame content by source, time.

Year	Reuters		<i>Defense One</i>		<i>Foreign Affairs</i>	
	Articles	Frames	Articles	Frames	Articles	Frames
2012	26	0	0	0	0	0
2013	11	0	0	0	2	0
2014	47	2	18	1	1	0
2015	114	17	23	7	6	0
2016	367	47	40	8	4	0
2017	766	56	77	10	4	0
2018	999	67	145	7	19	1
2019	1106	58	215	10	19	0

Appendix B: Annotation Guidelines

Guidelines for annotation that we applied for this report are available on CSET's [GitHub](#).

Endnotes

¹ Tom Simonite, "For Superpowers, Artificial Intelligence Fuels New Global Arms Race," *Wired*, August 8, 2017, <https://www.wired.com/story/for-superpowers-artificial-intelligence-fuels-new-global-arms-race/>.

² "China is Winning a New Global Arms Race," *Bloomberg Markets and Finance*, November 6, 2018, <https://www.youtube.com/watch?v=1Nr6daj4rBA>.

³ Julian E. Barnes and Josh Chin, "The New Arms Race in AI," *Wall Street Journal*, March 2, 2018, <https://www.wsj.com/articles/the-new-arms-race-in-ai-1520009261>.

⁴ Stuart Armstrong, Nick Bostrom, and Carl Shulman, "Racing to the Precipice: A Model of Artificial Intelligence Development," *AI & Society* 31 (2016): 201-206; Remco Zwetsloot, Helen Toner, and Jeffrey Ding, "Beyond the AI Arms Race America, China, and the Dangers of Zero-Sum Thinking," *Foreign Affairs*, November 16, 2018, <https://www.foreignaffairs.com/reviews/review-essay/2018-11-16/beyond-ai-arms-race>; Elsa B. Kania, "The Pursuit of AI Is More Than an Arms Race," *Defense One*, April 19, 2018, <https://www.defenseone.com/ideas/2018/04/pursuit-ai-more-arms-race/147579/>; Crispin Rovere, "Explained: The Artificial Intelligence Race is an Arms Race," *The National Interest*, February 2, 2020, <https://nationalinterest.org/blog/buzz/explained-artificial-intelligence-race-arms-race-119326>.

⁵ Heather M. Roff, "The Frame Problem: The AI 'Arms Race' Isn't One," *Bulletin of Atomic Scientists* 75, Issue 3 (2019): 2019; Tim Hwang and Alex Pascal, "Artificial Intelligence Isn't an Arms Race," *Foreign Policy*, December 11, 2019.

⁶ Andrew Imbrie and Elsa Kania, "AI Safety, Security, and Stability Among Great Powers: Options, Challenges, and Lessons Learned for Pragmatic Engagement," *Center for Security and Emerging Technology*, December 2019, <https://cset.georgetown.edu/wp-content/uploads/AI-Safety-Security-and-Stability-Among-the-Great-Powers.pdf>.

⁷ On the role of narratives and legitimation strategies in national security, see e.g., Ronald R. Krebs, *Narrative and the Making of US National Security* (Cambridge: Cambridge University Press, 2015); Stacie E. Goddard and Ronald R. Krebs, "Rhetoric, Legitimation, and Grand Strategy," *Security Studies* 24, Issue 1 (2015): 5-36.

⁸ Our 4,000 articles are English-language articles from three different news outlets; one global news outlet, Reuters, and two U.S. news outlets, *Defense One* and *Foreign Affairs*.

⁹ Radina Gigova, "Who Vladimir Putin Thinks Will Rule the World," *CNN*, September 2, 2017, <https://www.cnn.com/2017/09/01/world/putin-artificial-intelligence-will-rule-world/index.html>.

¹⁰ Elsa Kania and Rogier Creemers, "Xi Jinping Calls for 'Healthy Development' of AI (Translation)," *New America*, November 5, 2018, <https://www.newamerica.org/cybersecurity-initiative/digichina/blog/xi-jinping-calls-for-healthy-development-of-ai-translation/>.

¹¹ “Policy and Investment Recommendations for Trustworthy Artificial Intelligence,” High-Level Expert Group on AI, *European Commission*, June 26, 2019, <https://ec.europa.eu/digital-single-market/en/news/policy-and-investment-recommendations-trustworthy-artificial-intelligence>.

¹² Michael Kratsios, Remarks at “How the United States Can Maintain Its Lead in the Global AI Race,” *National Press Club*, September 10, 2019, <https://itif.org/events/2019/09/10/how-united-states-can-maintain-its-lead-global-ai-race>.

¹³ We retrieved articles from Reuters and *Foreign Affairs* on January 14, 2020, using Georgetown’s *Factiva* subscription in compliance with its terms of use. Our coverage of *Foreign Affairs* is restricted to the print journal. We collected *Defense One* content on February 6, 2020, from its [website](#).

¹⁴ See Appendix B for the annotation guidelines.

¹⁵ We refer to our search results as e.g. “AI articles” for brevity, but our query for mentions of “artificial intelligence” or “AI” may omit some articles about AI that did not include either term. Among the Reuters articles retrieved were a small number—fewer than 10—of science review-style articles that contained summary paragraphs referencing larger stories. As these reviews ran multiple days referring to the same stories, some of these summary paragraphs were repeated across reviews, and annotated repeatedly; however, readers may also have seen them repeatedly, increasing their relevance, so we considered them independent articles.

¹⁶ *Defense One* launched in July 2013.

¹⁷ As *Defense One* did not launch until July 2013, data for 2012 is only included for *Reuters*. See Kevin Baron, “Welcome to Defense One,” *Defense One*, July 15, 2013, <https://www.defenseone.com/news/2013/07/welcome-defenseone/66683/>.

¹⁸ As individuals and their affiliations were referred to with inconsistent naming conventions across articles, we performed manual entity disambiguation and aggregation.

¹⁹ Susan Decker, “Majority of Promising AI Startups Are Still Based in the U.S.,” *Bloomberg*, March 2, 2020, <https://www.bloomberg.com/news/articles/2020-03-03/majority-of-promising-ai-startups-are-still-based-in-the-u-s>.

²⁰ See, e.g., Ulrike Esther Franke, “Not Smart Enough: The Poverty of European Military Thinking on Artificial Intelligence,” *European Council on Foreign Relations*, December 2019.

²¹ As there are very few affiliations in frames with the critique purpose, it was not included.

²² Alex Krizhevsky, Ilya Sutskever, Geoffrey E. Hinton, “ImageNet Classification with Deep Convolutional Neural Networks,” <https://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf>.