KGC 2021 Workshop on Knowledge-Infused Learning

Using Contact, Content, and Context in Knowledge-Infused Learning: A Case Study of Non-Sequential Sales Processes in Sales Engagement Graphs

Hong Yung (Joey) Yip, Yong Liu, Amit Sheth

AI Institute, University of South Carolina, Columbia, SC, USA
Outreach Corporation, Seattle, WA, USA

Knowledge-Infused Learning (K-IL)

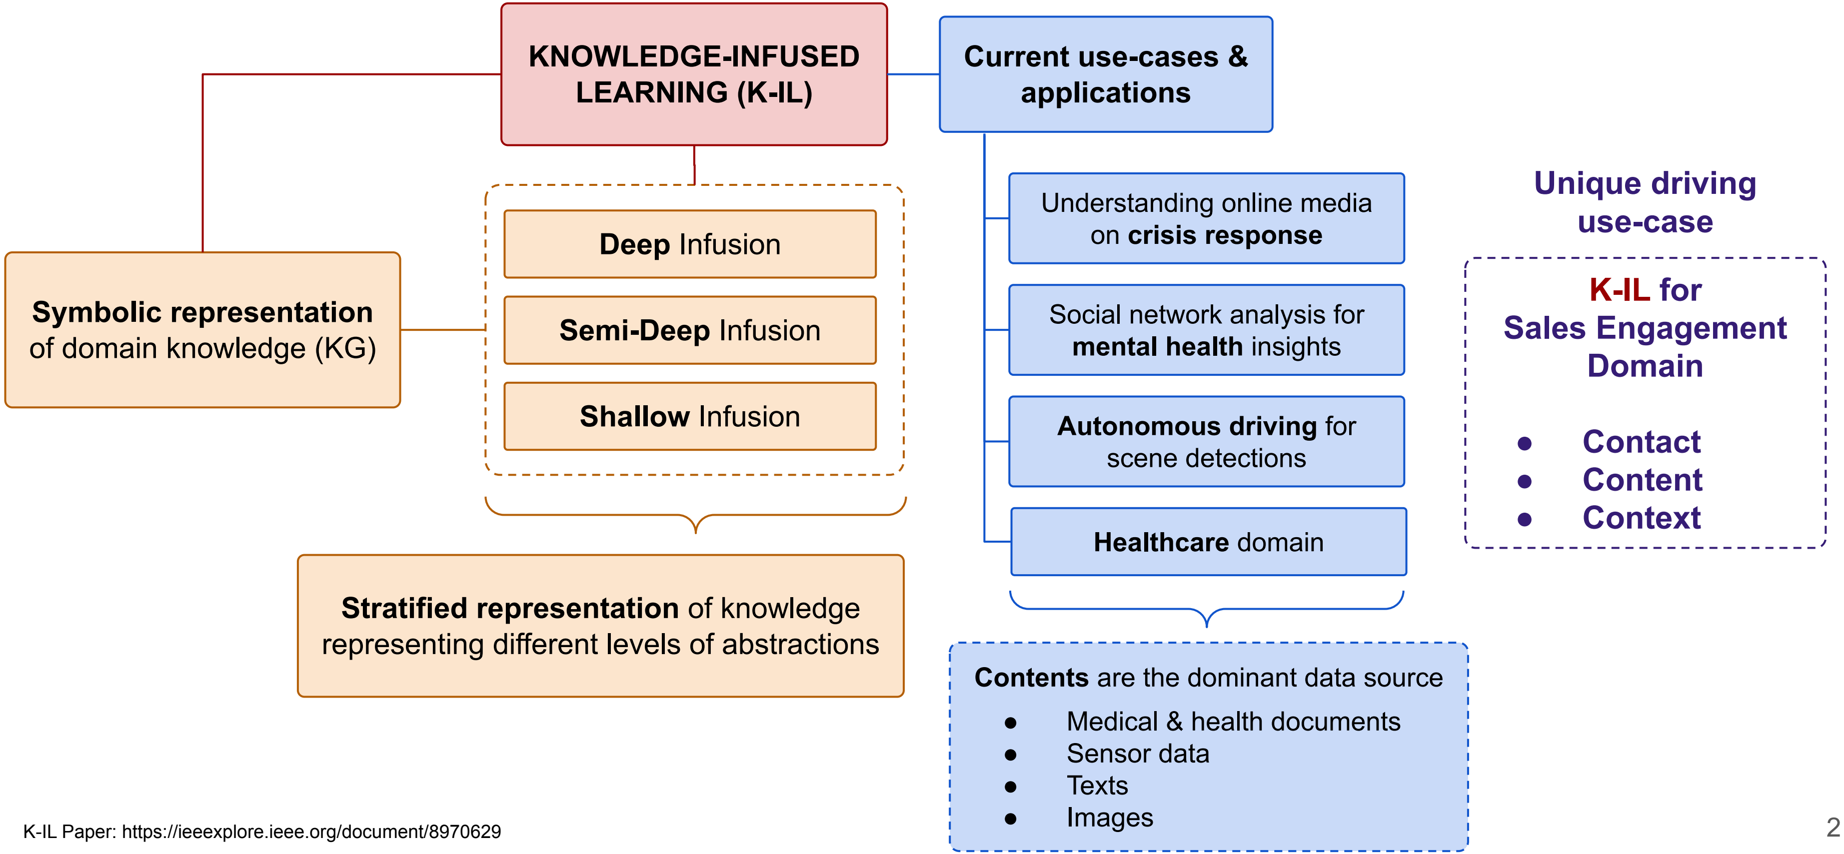

K-IL Paper: https://ieeexplore.ieee.org/document/8970629

Contact, Content, and Context (3Cs)

KNOWLEDGE-INFUSED LEARNING (K-IL)

KNOWLEDGE | LEARNING

Our proposal: Incorporate not just **contents**, but also **contacts** and **contexts** in the **K-IL approach**.

CONTACT

The knowledge about the **communication** and **buying preference** of the buyers and who needs to be talked to

CONTENT

The metadata and information of the past activities, emails, deal status & qualified leads' attributes.

Example: The **subject** and **body** of inbound **emails** provide information about **buyers intent** & relevant content for different sales activities.

CONTEXT

The **engagement history**, time-bounded deal-closing **constraints** and buyers **pain points**

Sales Engagement Domain

All 3Cs are **temporal** and **dynamic** in nature

Non-Sequential Sales Processes & Challenges

Multi-Activity

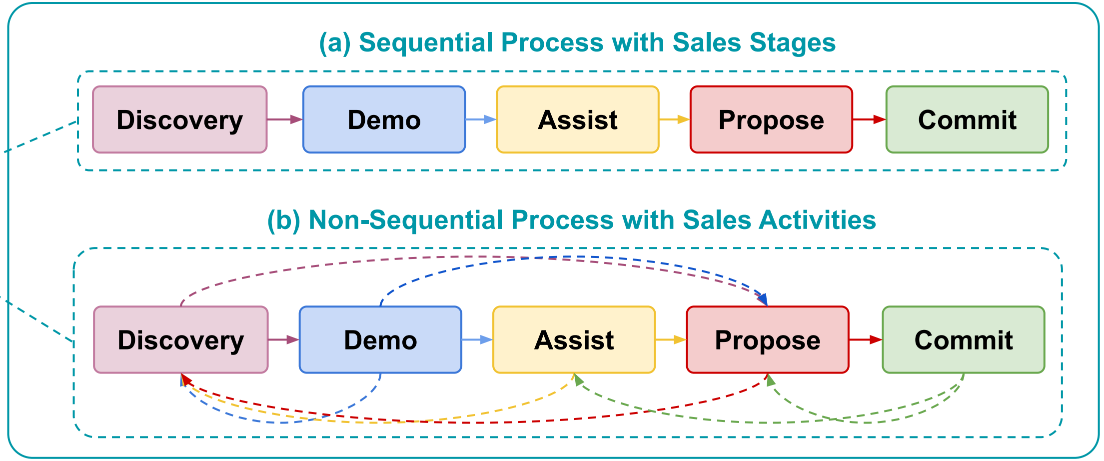

Sales Process

Multi-Actor

Seller

Sales Development Representative (SDR),
Customer Success Engineer,
Technical Solution Architect,
Account Executive (AE)

Buyer

Lead or contact or prospect,
Project champion,
Decision makers,
Budget holder

Challenges

1. Finding the right **contacts** to start or resume or accelerate a sale process
2. The lack of proper usage and understanding of **content** generated from both buyers and sellers
3. The insufficient **context** while transitioning and onboarding new accounts or closing a deal

Sales Engagement Platforms (SEPs)

A new category of software

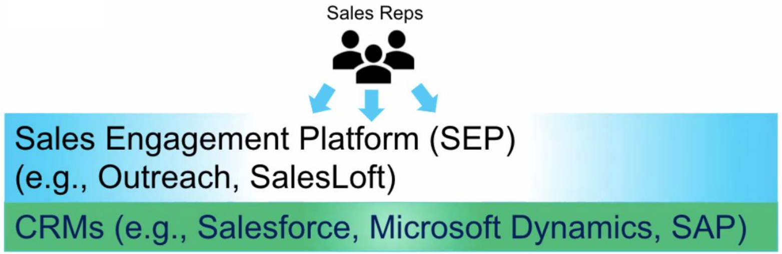

SEP

- SEP **encodes & automates** sales activities (sending emails, scheduling calls, meetings, etc) into workflows
- Enables sales reps to perform one-on-one personalized outreach **up to 10x**

Limitations

- **System of activities can happen at all sales stages.** Intents of email during discovery and commit are different.
- **Measuring the progress of active opportunities** and forecasting a deal outcome **based solely on engagement activity** and sales performance metrics have become **less effective**

Opportunity

Need to **go beyond gleaning the surface measures** and **look deeper into domain modeling (i.e.: 3Cs with KGs)** and **associated processes**

The Vision: Sales Engagement Graph Framework

Challenge: **Contact**, **Content**, and **Context** (3Cs) require proper modeling and data mining to derive actionable insights

Our Contributions

1. Methodology Development

A **three-phase Sales Engagement Graph (SEG)** framework for non-sequential sales processes

2. Application Area

Discuss **the role of 3Cs** in all three phases with challenges and solutions proposed

Present **initial implementation** of Phase 1 & **design** for Phase 2 and Phase 3

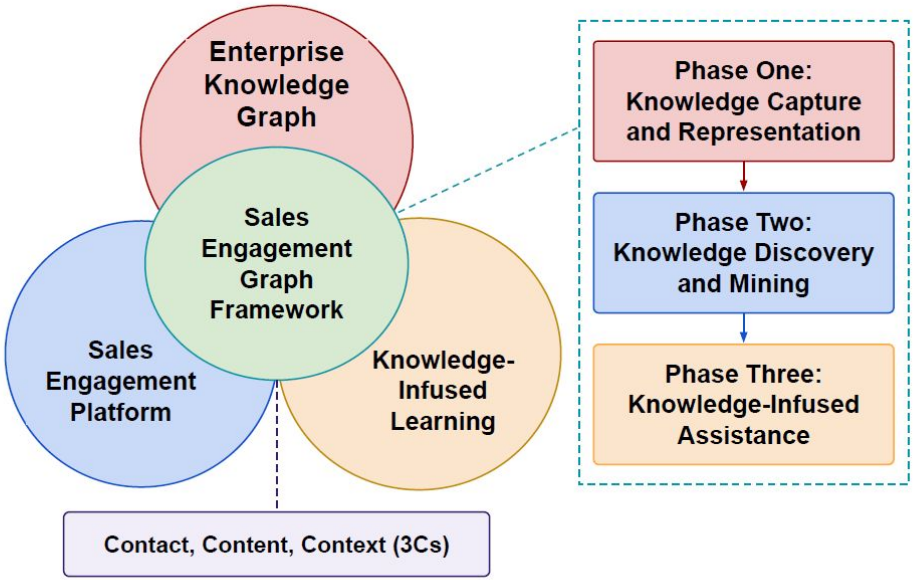

Data, Pipeline & Goals

Data: Emails, Meetings transcripts & notes, Engagement logs (email received & sent)

Three-phase Sales Engagement Graph (SEG)

Capture & unfold the underlying sales patterns, 3Cs & relationships between the types of sales activities to derive new insights

Support timely predictions and actionable recommendations for a given sales scenario

End goal: To have the **right context**, use the **best content** and **qualified contacts** to achieve **better guided engagement** between buyers & sellers to get a **positive engagement outcome** (meeting booked, deal closed etc.)

Phase One: Knowledge Capture and Representation

KNOWLEDGE CAPTURE

1. Determine use-cases
2. Scoping & documentation

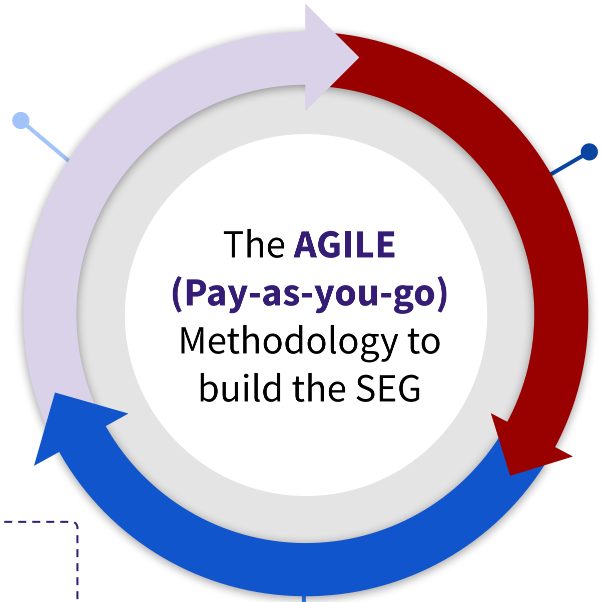

MULTI-SOURCE DATA INGESTION & INTEGRATION

3. Extract & integrate data (emails, engagement logs) while maintaining explicit provenance (with annotations of the data origin, creation time, and time of capture)

Primary Research Question: Can we capture the underlying actors involved in the sales process to surface the who-knows-who **(Contact)** information?

1. Who are the people in my organization has prior engagements with the buyer
2. Who are all the connected contacts for a given buyer

CONFLATION

4. Run entity disambiguation

Real-world person entity can have aliases, & abbreviations.
Unsupervised generative bayesian classifier with Expectation Maximization (EM).

Phase One: Knowledge Capture and Representation

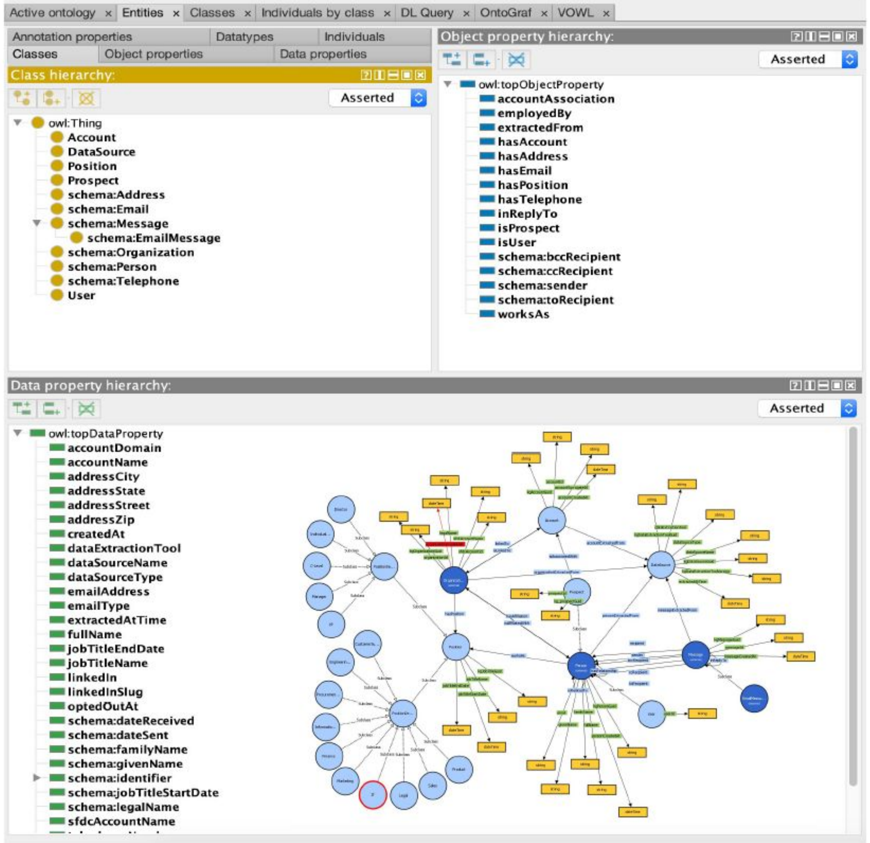

A preliminary version of the SEG ontology modeled around the **relationship** between the **actors** and **activities** in a sales process

Instantiated with people's information harvested from the **engagement logs** & **existing SEP's information** such as prospect's contact info.

Evaluation:

The pilot SEG derived from extracting over 4.7M engagement logs events (emails sent and received)

Harvested **64K newly discovered person** entities (**20% increase** in terms of number of people) - translate to a better coverage of all relevant people involved in the process

Conflation accuracy achieved 83% in terms of F1 score

Phase Two: Knowledge Discovery and Mining (Design)

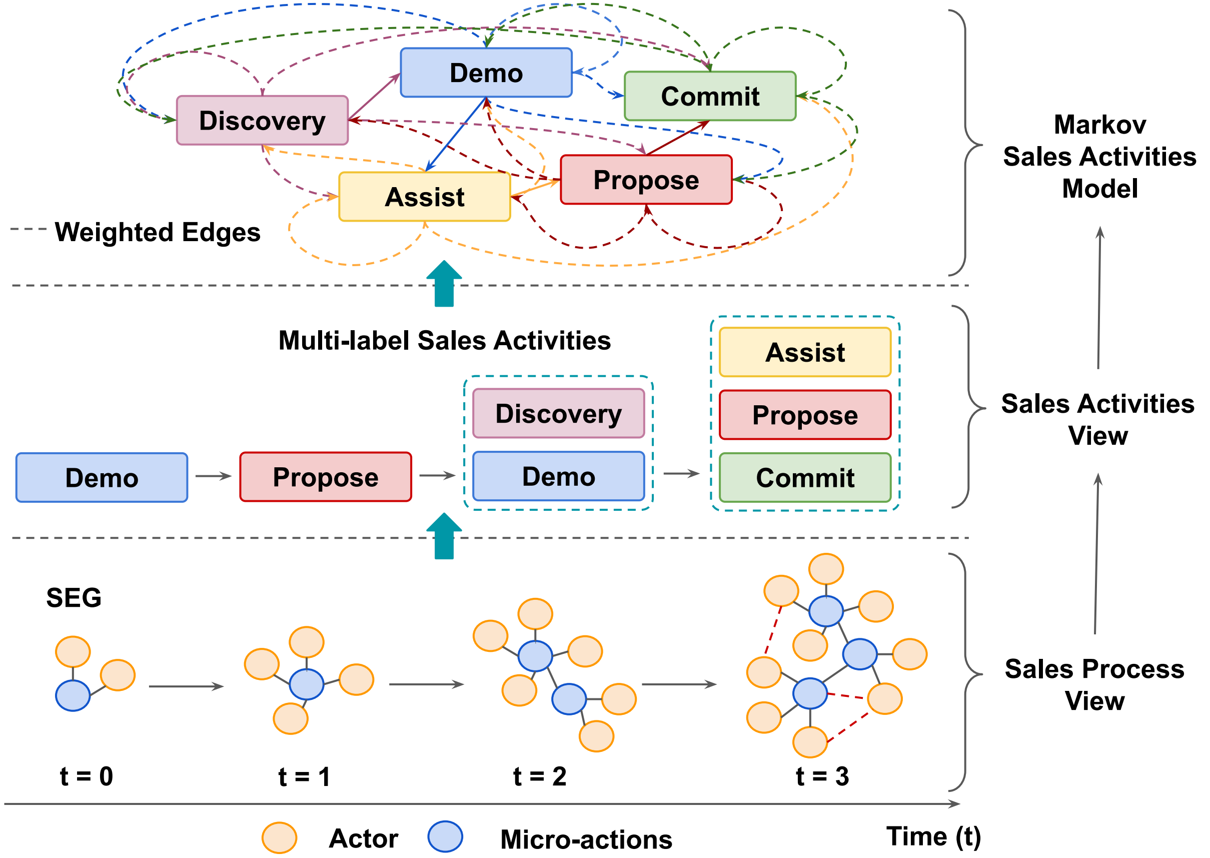

Goal Two: Establish a Probabilistic Markov Sales Activities Model: Model the relationships between the types of sales activities

Objective: Surface series of action paths to allow decision makers to analyze the sales activities that are often revisited & reshape sales objectives

Goal One: Discovery and Validation of Non Sequential Sales Processes by Multi-label Subgraph Classification: Labeling of higher-order semantics to temporal SEG interconnected 3Cs: meeting data, sender, recipients info, email signature, email intent

Primary Research Question: Can we unfold the dynamics between the types of sales activities within the time-bounded SEG?

Phase Three: Knowledge-Infused Assistance (Design)

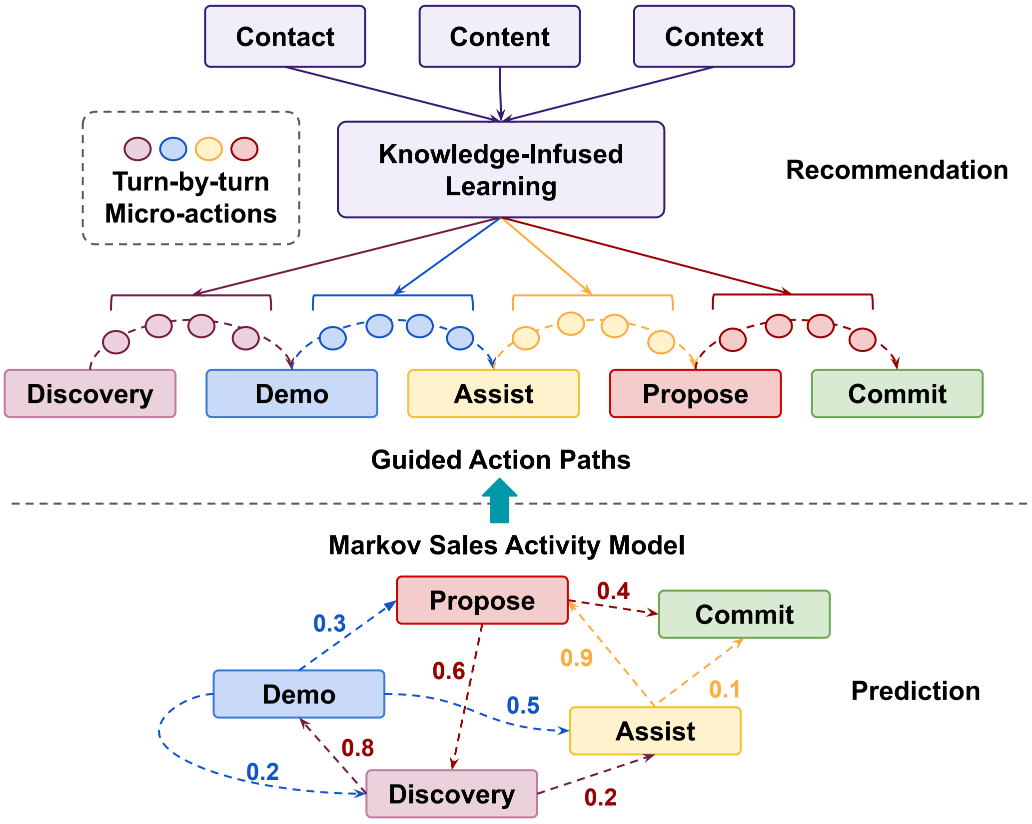

Goal Two: Recommending Turn-by-Turn Micro- Actions. Break the best action path down into series of micro-actions (eg. sending emails, setup meetings, etc)

Objective: Recommend turn-by-turn actions that produce the highest conversion rate

Goal One: Finding the Best Action Path.

Objective: Calculate impact & predict the outcome for an action path based on weighted transition probabilities of each sales activity

Primary Research Question: Can we recommend the next best action to maximize a deal's closing rate?

Example

BEFORE

Sales Reps: What should I do next to push for a positive engagement outcome or close the deal?

Phase 1: SEG captures the who-knows-who (Contact) information

Phase 2: With the Markov Sales Activities Model, we know the best action path(s) for this particular prospect based on historical dealings (how to proceed) (Content & Context)

Phase 3: Knowledge-Infused Assistance (3Cs)

Contact: Based on prior engagements, we know what is the communication preference (e.g., meeting call) and the best time to initiate the contact (e.g., morning hours)

Content: What materials, template, pricing, and incentives to use

Context: What was discussed and at which sales activities

AFTER

The next best action is to **schedule a meeting call between 8 and 10** in the morning with the **VP of sales** and **present a product demo**.

Conclusion

- The current COVID and post-COVID era have seen a shift in the sales practice of modern B2B (e.g., transitioning towards online-based): need to re-evaluate their strategic positions and offerings to the public.

- **Application:** While existing SEPs increase the overall productivity of sales reps, continuously providing "intelligence" needed for the users is the next evolutionary goal for SEPs.

- **Methodology:** The non-linear sales process provides a driving use case for us to envision a three-phase SEG framework that goes beyond incorporating not just **content**, but also **contacts** and **contexts** in the K-IL approach.

- We believe it is the step towards (a) a more holistic K-IL and (b) disrupting the next-era of the sales engagement and ultimately the customer engagement industry.

Acknowledgments: We would like to thank **Outreach Corporation** for supporting this research through an internship for the first author. We would also like to extend our acknowledgements, in no particular order, to Pavel Dmitriev, Balaji Shyamkumar, XinXin Sheng, Yong-Gang Cao, Sung Jeon, Rafi Anuar, Rajit Joseph, John Chu, Andrew Brooks, Mengyue Liang, Yifei Huang, Abhi Abhishek, and Andrew Herington for their continuous support in this project.

References

[1] Noy, N., Gao, Y., Jain, A., Narayanan, A., Patterson, A., Taylor, J. (2019). Industry-scale knowledge graphs: lessons and challenges. Communications of the ACM, 62(8), 36-43.
[2] Brandes, U., Gaertler, M., Wagner, D. (2003, September). Experiments on graph clustering algorithms. In European Symposium on Algorithms (pp. 568-579). Springer, Berlin, Heidelberg.
[3] Sheth, A., Gaur, M., Kursuncu, U., Wickramarachchi, R. (2019). Shades of knowledge-infused learning for enhancing deep learning. IEEE Internet Computing, 23(6), 54-63.
[4] Liu, Y., Dmitriev, P., Huang, Y., Brooks, A., Dong, L., Liang, M., ... Nguy, B. (2020). Transfer learning meets sales engagement email classification: Evaluation, analysis, and strategies. Concurrency and Computation: Practice and Experience, https://doi.org/10.1002/cpe.5759.
[5] Najafabadi, M. M., Villanustre, F., Khoshgoftaar, T. M., Seliya, N., Wald, R., Muharemagic, E. (2015). Deep learning applications and challenges in big data analytics. Journal of big data, 2(1), 1-21.
[6] Korhonen, J. J., Melleri, I., Hiekkanen, K., Helenius, M. (2014). Designing data governance structure: An organizational perspective. GSTF Journal on Computing (JoC), 2(4).
[7] Wilcke, X., Bloem, P., De Boer, V. (2017). The knowledge graph as the default data model for learning on heterogeneous knowledge. Data Science, 1(1-2), 39-57.
[8] Sheth, A. P. (2005). From semantic search integration to analytics. Semantic Interoperability and Integration, Dagstuhl Seminar Proceedings, http://drops.dagstuhl.de/opus/volltexte/2005/46
[9] Shi, B., Weninger, T. (2018, April). Open-world knowledge graph completion. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 32, No. 1).
[10] Gomez-Perez, J. M., Pan, J. Z., Vetere, G., Wu, H. (2017). Enterprise knowledge graph: An introduction. In Exploiting linked data and knowledge graphs in large organisations (pp. 1-14). Springer, Cham.
[11] Sashi, C. M. (2012). Customer engagement, buyer-seller relationships, and social media. Management decision. Vol. 50 No. 2, pp. 253-272. https://doi.org/10.1108/00251741211203551
[12] Ji, S., Pan, S., Cambria, E., Marttinen, P., Yu, P. S. (2020). A survey on knowledge graphs: Representation, acquisition and applications. arXiv preprint arXiv:2002.00388.
[13] Kazemi, S. M., Goel, R., Jain, K., Kobyzev, I., Sethi, A., Forsyth, P., Poupart, P. (2020). Representation Learning for Dynamic Graphs: A Survey. Journal of Machine Learning Research, 21(70), 1-73.
[14] Rossi, E., Chamberlain, B., Frasca, F., Eynard, D., Monti, F., Bronstein, M. (2020). Temporal graph networks for deep learning on dynamic graphs. arXiv preprint arXiv:2006.10637.
[15] Deng, S., Rangwala, H., Ning, Y. (2020, August). Dynamic Knowledge Graph based Multi-Event Forecasting. In Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery Data Mining (pp. 1585-1595).
[16] Ahmad, I., Akhtar, M. U., Noor, S., Shahnaz, A. (2020). Missing link prediction using common neighbor and centrality based parameterized algorithm. Scientific reports, 10(1), 1-9.
[17] Hamilton, W. L. (2020). Graph representation learning. Synthesis Lectures on Artificial Intelligence and Machine Learning, 14(3), 1-159.
[18] Xiong, W., Yu, M., Chang, S., Guo, X., Wang, W. Y. (2018). Oneshot relational learning for knowledge graphs. arXiv preprint arXiv:1808.09040.
[19] Motoda, H. (2007, June). Pattern discovery from graph structured data-a data mining perspective. In International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems (pp. 12-22). Springer, Berlin, Heidelberg.
[20] Fish, A. N. (2012). Knowledge automation: how to implement decision management in business processes (Vol. 595). John Wiley Sons.
[21] Alambo, A., Gaur, M., Thirunarayan, K. (2020). Depressive, drug abusive, or informative: Knowledge-aware study of news exposure during COVID-19 outbreak. arXiv preprint arXiv:2007.15209.
[22] Gaur, M., Alambo, A., Sain, J. P., Kursuncu, U., Thirunarayan, K., Kavuluru, R., ... Pathak, J. (2019, May). Knowledge-aware assessment of severity of suicide risk for early intervention. In The World Wide Web Conference (pp. 514-525).
[23] Wickramarachchi, R., Henson, C., Sheth, A. (2020). An evaluation of knowledge graph embeddings for autonomous driving data: Experience and practice. arXiv preprint arXiv:2003.00344.
[24] Rehman, S. U., Khan, A. U., Fong, S. (2012, August). Graph mining: A survey of graph mining techniques. In Seventh International Conference on Digital Information Management (ICDIM 2012) (pp. 88-92). IEEE.
[25] Sun, K., Lin, Z., Zhu, Z. (2020, April). Multi-stage self supervised learning for graph convolutional networks on graphs with few labeled nodes. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 34, No. 04, pp. 5892-5899).
[26] Eddy, S. R. (2004). What is a hidden Markov model?. Nature biotechnology, 22(10), 1315-1316.
[27] Gejingting, X., Ruiqiong, J., Wei, W., Libao, J., Zhenjun, Y. (2019, July). Correlation analysis and causal analysis in the era of big data. In IOP Conference Series: Materials Science and Engineering (Vol. 563, No. 4, p. 042032). IOP Publishing.
[28] Tandel, S. S., Jamadar, A., Dudugu, S. (2019, March). A survey on text mining techniques. In 2019 5th International Conference on Advanced Computing Communication Systems (ICACCS) (pp. 1022-1026). IEEE.
[29] Medina, M., Altschuler, M., Kosoglow, M. (2019). Sales Engagement: How The World's Fastest Growing Companies are Modernizing Sales Through Humanization at Scale. Hoboken, New Jersey : John Wiley Sons.
[30] Kursuncu, U., Gaur, M., Sheth, A. (2019). Knowledge infused learning (K-IL): Towards deep incorporation of knowledge in deep learning. arXiv preprint arXiv:1912.00512.