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

Hong Yung Yip[a], Yong Liu[b] and Amit Sheth[a]

[a]*Artificial Intelligence Institute, University of South Carolina, Columbia, SC, USA*
[b]*Outreach Corporation, Seattle, WA, USA*

Abstract

A sequential view of the sales processes no longer depicts the modern business-to-business (B2B) practice. While existing Sales Engagement Platforms (SEP) improve the overall productivity of sales representatives, we believe providing higher-level "intelligence" to support guided engagement between the buyers and sellers is the next evolutionary goal for SEPs. However, there exist challenges of surfacing and incorporating the sales engagement domain knowledge: **Contact**, **Content**, and **Context** (3Cs) into existing knowledge-infused learning approaches. In this paper, we describe our vision towards a Sales Engagement Graph (SEG) for non-sequential sales processes, propose a three-phase SEG framework that uses 3Cs: (1) Phase One: Knowledge Capture and Representation, (2) Phase Two: Knowledge Discovery and Mining, and (3) Phase Three: Knowledge-Infused Assistance, and share our experiences towards an initial implementation (phase one) and conceptualization and design (phase two and three) of the framework. In phase one, we demonstrate the implications of building a dynamic SEG with an agile and a use-case driven design process to surface the Contact information. In phase two, we illustrate the importance of abstraction for temporal pattern discovery and mining for contextual understanding. Lastly in phase three, we provide a straw-man approach to infuse the 3Cs with the Knowledge-Infused Learning approach for guided sales assistance and engagement, which goes beyond the common content-only knowledge-infused learning.

Keywords

Sales Engagement Graph, Sales Engagement Platform, Knowledge-Infused Learning, Contact, Content, Context

1. Introduction

Sales is a process between two or more parties in which an offering is exchanged with an equivalent value in return. A traditional view of a sales process runs through the five sales stages starting with (1) **discovery** followed by (2) **demo**, then (3) **assist**, (4) **propose**, and lastly (5) **commit** (Figure 1a). However, such sequential progression is not an accurate reflection of what actually happened in a modern business-to-business (B2B) purchase [1]. The major sales challenges are (a) finding the right **contacts** to start or resume or accelerate a sale process, (b) the lack of proper usage and understanding of **content** generated from both buyers and sellers, and (c) the insufficient **context** while transitioning and on-boarding new accounts or closing a deal. Hence, the conventional approach of determining an opportunity's probability of closing based on completion of prior sales stages (chronologically) is no longer sustainable. There is a need to re-examine the process in a non-sequential fashion (Figure 1b). A better depiction is to re-imagine the sales "stages" as "activities" as the term "stages" implies a form of linearity and sales representatives (sales reps) are more interested in the relationships between the five types of activities to surface insights to address the sales challenges.

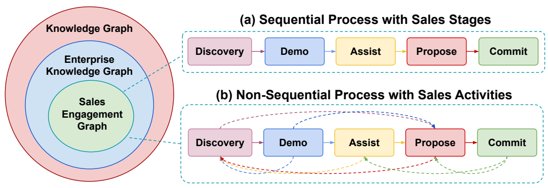

Figure 1: A Taxonomy of Knowledge Graph and Sales Process

Recent years have seen positive disruptions in the sales industry with Sales Engagement Platforms (SEP). Unlike a Customer Relationship Management (CRM) which maintains a system of records, an SEP captures a system of activities (scheduling meetings, sending automated emails, etc) and automates them as sequences [4]. While an SEP is great at supporting sales reps in their day-to-day tasks, it is still unclear how the system of activities affects the overall sales progression. Scheduling meetings and sending emails can happen at discovery, demo, and/or during commit activity. It is thus desirable to translate the system of activities into a higher-order view

The Knowledge Graph Conference 2021 Workshop on Knowledge-Infused Learning (K-IL)
hyip@email.sc.edu (H. Y. Yip); yong.liu@outreach.io (Y. Liu); amit@sc.edu (A. Sheth)

K-IL Workshop Proceedings (aiisc.ai/KiL2021)

[1]http://bit.ly/non-sequential-sales-processes

of the non-linearity sales process to suggest next actions or coaching needed by the sales reps [1].

With the growth of user-generated data from plethora of sales communication channels, organizations often find themselves having to sift through the sea of data to determine the next best action for a given sales objective. Although the state-of-the-art Artificial Intelligence (AI) technologies have seen promising applications [5], sales engagement intelligence requires approaches that take into account of the unique patterns and domain knowledge: **Contact**, **Content**, and **Context** (hereinafter referred to as **3Cs**) information which require proper modeling and data mining to derive actionable insights. Orthogonally, the common data storage practice with data silos suffers from the lack of semantic interoperability. With the current developments in knowledge graph technologies (KG) [2], it is possible to design an approach that represents common substrate of knowledge within an organization, supports integration of heterogeneous information, reasoning, and inference of new facts. However, having a semantically coherent and linked smart data only provide the foundation for shedding the lights of the complex sales process. There is a need to provide suggestions and guidance to help sales reps to move the needle in a more timely fashion.

Recent advances in Knowledge-Infused Learning (K-IL) [3] holds promises to infuse knowledge with statistical data-driven techniques. However, the current common knowledge infused methods are mostly focused on content (e.g., in health care domain, medical and health documents are the dominant content source for knowledge infused learning), thus it is this paper's purpose to go beyond the current prevailing approach and propose to incorporate not just content, but also contacts and contexts in a holistic knowledge infused learning. The non-linear sales process provides a driving use case for us to study how to move towards such a goal so that the methodologies presented in this paper can be generalized and applied to other domains when needed. To the best of our knowledge, this is the first paper to describe the use of 3Cs in knowledge infused learning and its application in sales engagement domain.

The contributions of this paper are thus the following:

1. A three-phase Sales Engagement Graph (SEG) framework for non-sequential sales processes is proposed: (a) Phase One: Knowledge Capture and Representation, (b) Phase Two: Knowledge Discovery and Mining, and (c) Phase Three: Knowledge Infused Assistance that synthesizes the KG, SEP, and K-IL domains and incorporates 3Cs for holistic knowledge infused learning (Figure 2).
2. The role of 3Cs (Contact, Content, and Context) in all three phases is discussed with challenges and solutions proposed.
3. Initial implementation of Phase 1 and design for Phase 2 and Phase 3 are presented

[2] https://2020-us.semantics.cc/knowledge-graphs-and-their-central-role-big-data-processing-past-present-and-future

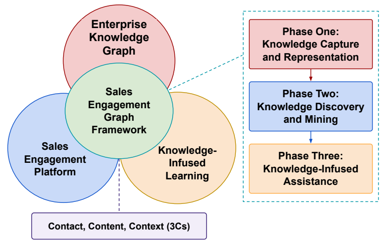

Figure 2: Sales Engagement Graph Framework

The paper is structured as follows. In section 2, we present the background and related work. In section 3, we describe our vision of the proposed SEG framework. We breakdown the framework into three distinct phases and provide a straw-man approach to conceptualize them (phase one is prototyped while phase two and three are on-going work). Lastly in section 4, we conclude the ideas presented in the paper.

2. Background and Related Work

Sales Engagement Domain. Sales is a multi-actor and multi-activity process. We define multi-actor as the players involved in a given sales process such as sellers (a Sales Development Representative (SDR), an Account Executive (AE), a Customer Success Engineer, and a Technical Solution Architect, etc) and buyers (a lead or contact or prospect, a budget holder, a project champion, and other key stakeholders and decision makers). Whereas for multi-activity, we define them as five types of sales activities: (1) **discovery** discusses the buyers' pain points and introduces the product capability, (2) **demo** involves demonstrating the product features to the buyers, (3) **assist** helps the buyers to formulate a plan based on their goals and budgets, (4) **proposes** provides relevant and targeted offers, and (5) **commit** focuses on overcoming the challenges of closing the deal (privacy, security, pricing concerns, etc) [1]. These activities are conventionally viewed as sales "stages" [1]. We define the individual action involved in these activities such as sending emails, making phone calls, scheduling meetings as **"micro-action"**. When micro-actions are executed in sequences, we define them as **"system of activities"**. However, with the evolution of internet technologies and the advent of social media platforms, the sales engagement domain is experiencing a paradigm shift in the sales process. As these platforms provide easy access and abundance of contextually relevant information to the buyers during their due diligence, there is a noticeable reduction in offline engagement between buyers and sellers in the current online-and-social-media-based economy [3]. It is very likely that the prospects will reach out to the seller and the first engagement starts directly with the **demo** instead of **discovery** activity (more on Section 3.2).

Sales Engagement Platform (SEP). An SEP turns a company's sales process into sequences and executes them in an orderly fashion. It streamlines the sales process through integrations with existing workflows across various sales tools (CRM, Content Management System (CMS), etc) and digital communication channels [4, 29]. Most SEPs (e.g., Outreach [4], SalesLoft [5]) in the current market are workflow-based automation tools where they automate and parallelize the system activities for a wider outreach and overall improved sales productivity. Despite their sophistication, the growth of an SEP's users and sales activities generated data has propelled the need to move from **workflow automation** to **knowledge automation** [6] [7] [20]. We define knowledge automation as the automatic capturing and digitizing of entities and relationships to form semantically enriched "knowledge items" that can be reused and repurposed to solve business problems. Furthermore, measuring the progress of active opportunities and forecasting a deal outcome based solely on engagement activity and sales performance metrics have become less effective [1]. We need to go beyond gleaning the surface measures and look deeper into modeling the sales engagement domain and associated processes, to which we identify the three important domain knowledge for a positive engagement outcome:

- **Contact:** The knowledge about the communication and buying preference of the buyers and who needs to be talked to.
- **Content:** The metadata and content [28] of the past activities, correspondence emails, status of the deal and referrals, relevant phone numbers and qualified leads' attributes. For example, the subject and body of inbound emails provide contextual information about buyers intent [4], while the email templates and meeting notes provide relevant content used in different sales activities.
- **Context:** The engagement history, time-bounded deal-closing constraints and buyers pain points

Knowledge Graph (KG). Inspired by human problem solving and reasoning, a KG is a form of semantic web technologies that represents knowledge (commonly formalized as an ontology in the RDF/OWL format in accordance with the W3C standard) for systems to gain the "intelligence" to solve complex tasks [1]. It holds promises to be the default data model for next-generation AI systems [7] that are explainable and interpretable, grounded with human knowledge [9]. Contrary to the semantic web-scale KG, the sales engagement domain operates under the *Semi-Closed-World Assumption* which contains a finite set of actors and activities [11], but still leverages some form of the external information such as people network in LinkedIn. An Enterprise Knowledge Graph (EKG) is a subset of the open-domain KG and is often tied to business processes [10]. A well-modeled EKG is posited to provide (a) ontological committed knowledge within an organization, thereby facilitating knowledge reusability across different facets of organizational units, products, and applications, and (b) allow the integration of heterogeneous information and external ontologies [8]. We believe the next step in unlocking the value of represented knowledge is to assimilate them to uncover hidden patterns and possibly surface the non-linear relationships between the types of sales activities.

Knowledge-Infused Learning (K-IL). Data mining on KG has become increasingly popular due to the advent of graph representation learning techniques [17] and their demonstrated strengths in learning complex and relational data [12]. Some of their applications include pattern discovery from graph-structured data [19], representation learning for dynamic graphs [13], temporal graph mining on sequences of timed events [14], multi-event forecasting [15], missing link predictions [16], and KG completion [18]. While they excel with adequate and high-quality labeled datasets, most are limited in their ability to perform higher level tasks that involve generality and explainability [3]. Knowledge-Infused Learning (K-IL) exploits the symbolic representation of domain knowledge and infuses them with various degrees (shallow, semi, and deep) and at different levels of abstraction in the latent layers of statistical data-driven techniques [3]. Some of the K-IL applications include understanding online media on crisis response [21], social network analysis for mental health insights [22], and in the autonomous driving for scene detections [23]. We analogize sales engagement to driving and that we need "Global Positioning System (GPS)"-like guidance. Hence, we posit that by fusing a well-modeled EKG with K-IL, we will be able to unfold the underlying sales patterns (i.e., 3Cs) to derive new insights and generate explainable recommendation to guide the next best action with traceability.

[3] https://www.gartner.com/en/newsroom/press-releases/2019-11-13-gartner-says-the-traditional-linear-sales-process-no-

[4] http://outreach.io/

[5] https://salesloft.com/

[6] https://www.arago.co/knowledge-automation/

[7] https://shelf.io/blog/2021-outlook-ai-knowledge-automation-and-the-evolving-role-of-remote-agents/

3. Sales Engagement Graph (SEG) Framework

To move towards our vision of a SEG for non-sequential sales processes, we propose a **three-phase SEG framework** (Figure 2), motivated by (1) the need to represent the sales engagement domain (phase one), (2) our interest in unfolding the underlying relationships between the types of sales activities (phase two), and (3) for the SEG to ultimately support timely predictions and actionable recommendations for a given sales scenario, and integration with an existing SEP to achieve better guided engagement between buyers and sellers (phase three).

3.1. Phase One: Knowledge Capture and Representation

This phase is to lay the foundation of a proof-of-concept SEG. An example use-case is to **find the right actors to initiate or resume a sales engagement**. Insufficient people **context** while transitioning and onboarding new accounts can be discouraging even for seasoned SDRs [4]. Hence, our primary research question is: "Can we capture the underlying actors involved in the sales process to surface the who-knows-who **(Contact)** information?". Concretely, these are the competency questions that we want to address:

1. Who are the actors in my organization has prior engagements with a target actor
2. Who are the mutual acquaintances of an actor
3. Who are all the connected contacts for a given actor
4. The average response time per engagement

We believe that, aside from people's information in CRM systems, there is rich information in the engagement logs **(Content)**. An engagement log is produced by the SEP to record micro-actions and the metadata information such as names, email addresses, time and engagement activity, in- and outbound communications and roles related to the person of interest. Our process of building the SEG consists of a three iterative steps: (1) knowledge representation, (2) multi-source data ingestion and integration, and (3) conflation.

Knowledge Representation. We believe that the most productive use of a KG is not to be confined by the academic notions (e.g., adoptions of standardized markup schemes) nor fixated on the creation of a centralized graph in a waterfall fashion. Instead, we share similar ideologies with diffbot [8] that KG technologies (EKG in particular) should be viewed as a way to introduce knowledge automation by transforming mundane, high frequency

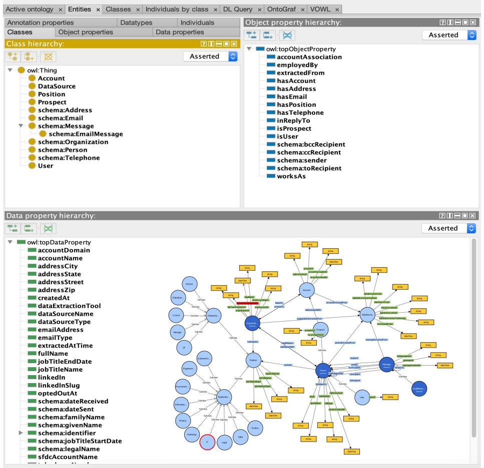

Figure 3: A Formal Definition of the SEG Ontology (snapshot)

and volume information into machine-readable representation, and automating knowledge workflow, which includes adding semantic and business meaning to data that are captured by an SEP. Unlike a semantic web-scale KG, most of the sales engagement domain knowledge are driven by the underlying data instances collected and extracted from enterprise data warehouses and CRM record systems. Hence, we adopt an agile (in a pay-as-you-go, modular, and iterative fashion) and a use-case driven approach in our knowledge modeling process, instead of the waterfall fashion of constricting the design with a list of pre-defined domain classes and relationships. We believe the implications of such design process (agile and use-case driven) allow us to enhance and grow the SEG graph organically, both at the schema and instantiation layer due to the dynamicity nature of the sales domain as well as to account for ever-changing requirements (e.g., updating the schema to absorb newly discovered knowledge items). A preliminary version of the SEG ontology modeled around the relationship between the actors and micro-actions in a sales process is illustrated in Figure 3 (depicted in Protege[9]). It is then instantiated with people's information harvested from the engagement logs enriched/integrated with an SEP's application contextual information such as prospect's contact info that are already existing in other data silos. The SEG is then used to drive and answer our research questions.

Multi-source Data Ingestion and Integration. A SEG can serve as an abstracted semantic layer to facilitate meaningful mapping (data viewed as entities rather than strings) between domain entities, however it is desirable that it should not be another silo. Instead, it should serve as a link to standardize and orchestrate interdisci-

[8]https://blog.diffbot.com/from-knowledge-graphs-to-knowledge-workflows/

[9]https://protege.stanford.edu/

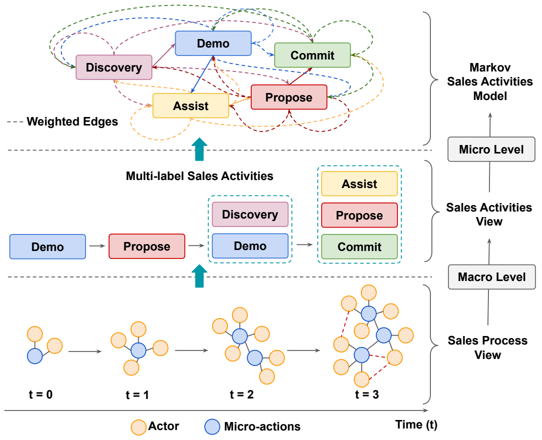

Figure 4: The Abstraction from a Multi-actor and Multi-activity Sales Process View to a Sales Activities View to a Markov Sales Activity Model

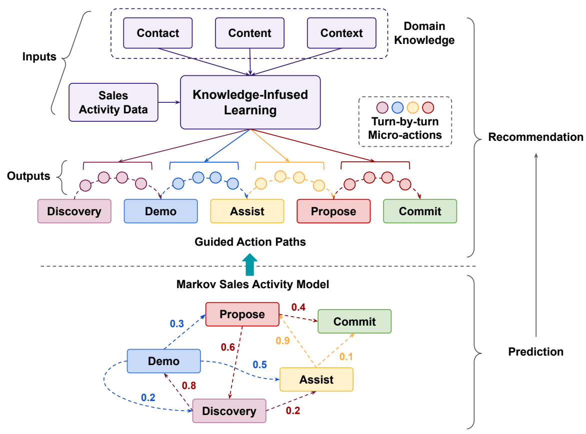

Figure 5: From Markov Sales Activity Model to Guided Action Paths with Knowledge-Infused Learning

plinary communication and information flows between isolated units in an organization. In our use-case, our data sources are (a) **enterprise data warehouses** which contain customer and organization data and (b) **engagement logs**. Maintaining explicit provenance is crucial during integration especially when making determinations in situations where information from both sources appear contradictory (e.g., mismatched email addresses). We made effort to preserve all relevant provenance metadata with annotations of the data origin, creation time, and time of capture to ensure data recency and relevancy.

Conflation. Conflation refers to the discovery, identification, and collation or merging of two or more (often duplicate) entries and instances referring to the same real-world entity under a single unique and consistent entity in the SEG. This step is particularly important because a real-world person entity can have aliases or signed-up under different names, abbreviations, and middle initials. Each person entry is assigned with a unique identifier in different silos. It is difficult to update the information across all entries in various silos without a conflated entity to be the key global identifier. In addition, a real-world person can play a different role in different organizations (a person can be a manager in company A and a customer of company B). We use the unsupervised generative bayesian classifier with Expectation Maximization (EM) algorithm for entity matching [10]. Some of the features include full name, address city, email address, account name, LinkedIn ID, and job title. In our experience, an email address is not a globally distinct identifier for a real-world person entity.

Evaluation of the Phase-1 Impact. Phase one saw a pilot implementation on one of the leading SEPs (Outreach [4]) with three organization data used. The resulting SEG is derived from extracting over 4.7M engagement logs events (emails sent and received). The pilot SEG contains approximately 300K unique person entities out of which 242K were unique prospects with an additional 64K newly discovered person entities (not presented in the existing data warehouse), which translated into a 20% increase in terms of number of people. Conflation accuracy achieved 83% in terms of F1 score where the F1 score is a harmonic mean of precision and recall. Answers to the list of competency questions (Section 3.1) were also provided.

3.2. Phase Two: Knowledge Discovery and Mining

The ideas described in this section draw inspirations from Jacco's framework [11] and they delineate our initial steps towards **discovering** and **mining** the underlying pattern of a sales process, in the form of **action path** from the SEG. We define an action path as a path of sales activities (e.g., discovery->demo->propose->commit that a sales process took to close the deal). Our primary research question is "Can we unfold the dynamics between the types of sales activities (in a per deal basis) at both macro (the full-length action path) and micro level (the transition from one sales activity to another) within a time-bounded SEG?", to which we further divide it into two sub-goals.

Goal One: Discovery and Validation of Non Sequential Sales Processes. We hypothesis that at the macro level, a sales process can be modeled as a directed eulerian temporal graph with a starting and an ending

[10] https://github.com/moj-analytical-services/splink

[11] https://winningbydesign.com/

activity. On the other hand, at the micro level, the transition from one sales activity to another is a non-sequential process. In a typical (mock-up) sales process scenario (Figure 4), suppose at t=0, the sales process (at the **sales process view**) starts with a micro-action: a sales rep initiates with an email to prospects. Then at t=1, the prospects involve more parties in the email conversation. Later at t=2, the sales rep schedules a follow-up meeting calls involving more parties. Lastly at t=3, decision makers are involved to commit the deal. While most SEPs are capable of capturing the system of activities (sequences of micro-actions), it is a major challenge to provide the insights and relevant **contexts** needed for next actions as there are no higher-order semantics attached (labeling of the five types of sales activities) to them. These micro-actions can happen at any timestep. We are interested in understanding the dynamics at each timestep, to which we put forward a need to restructure the sales process view to an **abstracted sales activities view** by labeling the SEG at each timestep at a given time granularity (minutes, hours, days, etc), formulated as a graph-labeling [24, 25] and clustering tasks [2]. At each timestep of various granularity, there can be occurrence of multiple sales activities due to the inherent nature of the sales process. At a coarse level labeling, we can identify signals from the metadata of the micro-action such organizer, attendees, start time, end start, meeting title, and agenda from meeting data, and sender, recipients from email messages. We can also identify a person's job title via the email signature. Job titles can serve as important contextual cues: engagement of more senior roles may indicate **propose** or **commit** activities. Whereas on a finer level, we can learn the intent of each micro-action (phase three) via the meeting description or email subject and body text [4]. Once the sales activities view is constructed, we can then validate the **non-sequential** and cyclical characteristics of the sales process.

Goal Two: Establish a Probabilistic Markov Sales Activity Model. The sales activities view can be further translated into a **model** with their transition from one to another quantified with weighted edges. With a Markov model [26], we can model the relationships between the types of sales activities in terms of their **closeness** (which activity is closely related to one another) and **activeness** (which activity co-occurs frequently with one another). For example, in this action path: discovery->demo->discovery, the discovery activities in the former and latter timestep carry different meanings albeit the same label, to which we can then apply possible causal and correlation analysis [27]. A series of **action paths** can be surfaced to support the bodywork for a contextualized and real-time guided engagement in phase three.

The implication of this phase is essentially to allow decision makers to analyze the sales activities that are often revisited in the sales process, and consequently reshape sales objectives around and redirect more effort and resources on such activities. This phase will be evaluated against ground truth data of prior closed dealings with each timestep (of defined granularity) classified and labeled by human annotators.

3.3. Phase Three: Knowledge-Infused Assistance

We propose our goals in this phase with the following research question in mind: "Can we **recommend** the next best action to maximize a deal's closing rate?", to which we further divide it into two sub-goals.

Goal One: Finding the Best Action Path. An organization's sales pattern can be modeled given a collective set of historical dealings. Figure 5 (bottom view) illustrates a (mock-up) Markov model for an organization. With the weighted transition probabilities of each sales activity known, we can then surface a series of action paths, calculate their impact (e.g., multiplication of transition probability from one sales activity to another), predict the outcome of a sales process for each action path, and hence determine the best action path. For example in Figure 5, the best action path is discovery->demo->assist->propose->commit.

Goal Two: Recommending Turn-by-Turn Micro-Actions. Although we have identified the best action path, it is still driven by many micro-actions, but not every micro-action is equally effective and will lead to a desirable outcome. Therefore, we are interested in breaking the action path down into series of micro-actions. We want to learn what micro-actions produce the highest conversion rate and subsequently recommend the series and/or combinations of micro-actions that best realize the corresponding path. Given we have the prior domain knowledge (3Cs) and an understanding on the nature of sales activity data (in a statistical manner), they can be inputs or state information to the K-IL approach (upper view of Figure 5) at various layers of infusions:

- **Shallow:** We can exploit the properties of the SEG: (a) the structure of the relationships between entities/ nodes, and (b) the semantic of each entity/node (e.g., a person node can have an associated job title node property/ attribute) and translate them into knowledge graph embeddings [12] as a form of shallow infusion in the K-IL approach.
- **Semi- and Deep:** A deeper contextual modeling [30] of the domain knowledge is infused at the semi- and deep infusion layers [3]. For example, from the **Contact** domain knowledge (based on prior engagements), we know what is the buyer's preferred mode of contact (e.g., meeting call) and the best time to initiate the contact (e.g., morning

hours). With infusion of such knowledge, the best micro-action to recommend to the sales rep is to schedule a meeting call between 8 and 10 in the morning with the buyer.

As micro-actions have direct correlation and impact on the sales outcome, the novelty and implication of this phase is to generate turn-by-turn (GPS-like guidance) micro-actions for a given sales scenario in a contextual and personalized fashion based on the 3Cs to **assist** the sales reps. They can then be integrated with an existing SEP to achieve better guided engagement between buyers and sellers. The evaluation of this phase on the system-recommended next best actions will be reviewed by the domain experts. While the work to implement this vision is still in its infancy, we believe this will bring the next round of innovation in SEPs.

4. Conclusion

A sales process ultimatum is to close as many deals as possible. The current COVID and post-COVID era have seen a shift in the sales practice of modern B2B (e.g., transitioning towards online-based), which prompted a need to re-evaluate their strategic positions and offerings to the public [12]. While existing SEPs increase the overall productivity of sales reps, continuously providing "intelligence" needed for the users to have a sustainable flywheel is the next evolutionary goal for SEPs. Without an understanding of the underlying sales pattern with the associated domain knowledge (3Cs) and the relationship between the types of sales activities, it is challenging to derive insights and deliver the relevant next best actions needed by sales reps. In this paper, we presented a three-phase SEG framework for non-sequential sales processes that goes beyond incorporating not just content, but also contacts and contexts in the K-IL approach. We believe it is the step towards 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.

[12]https://www.edelman.com/research/how-b2b-companies-are-adapting-covid

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). One-shot 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.