Spatial Predictive Queries

Mohamed Ali[1] Abdeltawab M. Hendawi[1,2]

[1]Center for Data Science, Institute of Technology, University of Washington Tacoma, WA, USA

[1]{mhali,hendawi }@uw.edu

[2]Department of Computer Science and Engineering, University of Minnesota, Minneapolis, MN, USA

[2]{hendawi}@cs.umn.edu

Duration: 1.5 Hours

***Abstract*—In this seminar, we address spatial predictive queries both in Euclidian spaces and over road networks. We provide a definition for various types of spatial predictive queries, describe current research trends, and envision future directions. We present practical application scenarios and emphasize the roadblocks that are holding industry back from the commercialization of spatial predictive queries. This seminar targets audience in mobile data management, spatiotemporal query processing, mobile crowd sourcing, and tracking of moving objects.**

I. TUTORIAL OUTLINE

The next generation of location based services would offer and recommend services for users according to their current locations as well as their future destinations [20], [21], [23], [25], [26], [29], [34], [37], [38], thanks to the widespread of mobile devices [17]. Nowadays, search engines, e.g., Microsoft Bing and Google, offer satisfactory location-aware search results that are based on the user's current location. However, these search engines remain in an infancy phase developing search techniques that takes into consideration the user's future (or intended) destination. Experience tells that the advertisement that targets the user's current location (e.g., a coupon in a nearby shopping mall) is not very effective. More specifically, it is too late to attract the user's attention to a nearby service because most users, by then, are already heading to a pre-planned destination. Industry believes that targeting the user with services that are around his future destination is more valuable to the user (from a relevance perspective) and more profitable to industry (from a market share perspective). However, experience tells that users are either reluctant to share their future destinations with search engines or are unaware of the value they may get by doing so. Consequently, the prediction of the user's future location combined with spatial query processing has been gaining tremendous interest in both the research and industrial communities. Entering this era of "future-location-aware" search engines requires the ability to process spatial predictive queries.

In this seminar, we survey the existing research and envision the future of spatial predictive query processing and optimization. The seminar is organized in the following key sections:

A. Part 1: Spatial Predictive Queries, What and Why?

In this part, we provide basic definitions for different types of predictive queries. Then, we show the importance of this topic through real-world example applications and systems. In general, predictive queries [19], [21], [25], [26] aim at answering inquiries about the anticipated future locations of a set of moving objects, either in an Euclidean space or over a road network.

The fundamental types of spatial predictive queries include: (a) *predictive point* query [19], [18], that finds out the objects that are most likely to show up around a specific location point in the space within a future time window, (b) *predictive range* query [26], [49], [60], where a user defines a query region rather than an exact point and asks for the list of objects expected to be inside the boundaries of that region within a specified future time window, (c) *predictive KNN* query [4], [43], [60], that gets the most likely K objects expected to be around a location of interest within a certain time period, (d) *predictive reverse-nearest-neighbor query* [4], [56], that finds out the objects expected to have the query region as their nearest neighbor. This query is useful in service distribution applications such as ad-hoc networking to assign mobile devices to the nearest communication service point, (e) *predictive aggregate* query [2], [20], [48], that predicts the total number of objects to appear at a desired location in a given next time range, and (f) *continuous predictive query* [28], [34], [37], [55], that allows any predictive query of the aforementioned query types to be stored at the server side and to be continuously reevaluated over a data stream of input locations through out its life in the system.

Spatial predictive queries can be utilized in a wide range of location-based services. Examples of these services include:

Weather Forecast. By considering the weather elements, e.g., clouds and tornadoes, as moving objects, the weather forecast applications can issue predictive range queries to predict the weather conditions at the areas of interest. In addition, smart warnings can identify those stationary/moving objects on the ground, e.g., homes/vehicles, that might be hit by a certain server weather event [35]. Accordingly, evacuation plans can be dispatched early before the occurrence of such dangerous events.

Ride Sharing Systems. Ride sharing services aim at linking the rider to the nearest driver and vice versa. These services can be improved by allowing the rider to submit a *predictive range* query. This query finds out the drivers that are most likely to show up around the rider's current location in the next few minutes. This approach would assist riders to plan their trips more efficiently and to reduce their waiting times.

Location-Aware Advertising. A store in a sale season executes a *predictive KNN* query to send electronic coupons to the K, e.g., fifty, costumers that most likely to show up around its location within the next t time units, e.g., next 30 minutes. This paradigm allows location-based advertising to go beyond nearby customers in the present time and to target possible nearby ones in the future. Sending coupons and promotions to these prospective customers would encourage them to plan ahead and stop by the store which, in turns, increases the effectiveness of advertising for both business owners and consumers.

Traffic Management Systems and Routing Services. *Predictive queries* in traffic management systems would improve traffic prediction results. Both Bing Maps and Google Maps offer routing services that calculate the travel time with and without traffic. In a typical situation, online traffic services announce a traffic jam after the fact, that is, after all interested users are either in the traffic jam itself or inevitably heading to the traffic jam. Another major problem in calculating the travel time for long trips is that the traffic pattern changes significantly while the driver is half-way in his trip. Predictive aggregate queries estimate the number of vehicles expected to be inside a certain area, e.g., down town, in the coming time interval, e.g., 30 minutes. Consequently, the estimated travel time of a trip would be enhanced by considering both the current and predicted future traffic conditions. Also, executing predictive queries against the road map before acquiring the route would eliminate the portions of the map that might be influenced by bad conditions, e.g., congested traffic, severe weather event. Accordingly, a safe routing service would avoid these undesirable areas.

B. Part 2: Research Trends

In this part, we survey the current research trends that form the field of spatial predictive queries. These trends can be summarized as follows:

(1) Query evaluation and optimization, in which the main concern is to find the optimal or at least a good enough strategy for executing the predictive queries [11], [10], [12], [16], [53], [52], [55].

(2) Prediction functions, which refer to the underlying prediction model employed to anticipate the next destination or the complete forthcoming trajectory of a given moving object. We talk about three categories of prediction models; (i) *Linearity-based prediction* [4], [43], [47], [50], [52], where the underlying prediction function is based on a simple assumption that objects move in a linear function in time along the input velocity and direction, (ii) *Historical-based prediction* [6], [14], [26], [29], [32], [31], [48], where the predication function uses object historical trajectories to predict the object's next location, and (iii) *complex prediction* [25], [49], [59], [61], where more complicated prediction functions are employed to realize better prediction accuracy.

(3) Spatio-temporal indexing techniques, which attempt to find an efficient way to store and retrieve moving objects data. We cover four categories of indices that are widely used within the context of processing spatial predictive queries. Each category is based on and derives its variants from the following basic structures; (i) *R-tree* [3], [44], [45], [46], [51], (ii) *B-tree* [8], [24], [58], (iii) *kd-tree* [5], [13], [39], [54], and (iv) *Quad-tree* [7], [39], [36], [41].

(4) Location uncertainty, which deals with the imprecise knowledge of the objects' locations, velocities, and directions. This trend tackles inference techniques that obtain the anticipated objects' locations given uncertain motion patterns and non-deterministic input values [42], [49], [59].

C. Part 3: Euclidean Space Versus Road Network

In this part, we address two different settings for spatial predictive queries. These queries can run in an Euclidean space or against a road network graph. The main difference between the Euclidean space and the road network is that the objects in the former are free to move anywhere in the given space without constrains. However, in the later, the objects' movements are constrained by the underlying road segments, intersections, and speed limits on each road. Also, in the Euclidean space, the Euclidean distance is the measure of distance between two locations. In a road network, the travel distance (or travel time) on the road segments that form the object's trip is considered to be the distance measure.

As an example system that handles spatial predictive queries on Euclidean space, we manifest the *Panda* framework [20]. We overview its data structures, prediction model, and key ideas behinds its efficiency and generality. As an example system that handles predictive queries on road networks, we present the *iRoad* [18] framework, its *predictive-tree (P-tree)* data structure [19], and associated query processing algorithms.

D. Part 4: Open Problems and Future Directions

In this part, we highlight a couple of challenges that are holding industry back from moving forward aggressively in pursuing spatial predictive query efforts: (i) Privacy, which aims at protecting the users' private information (e.g., exact locations, motion patterns, and historical trajectories) from being released while providing the users with meaningful services based on their location. In this seminar, we discuss several merits and drawbacks of each of the well-known techniques that achieve privacy in spatial non-predictive queries [1], [15], [33], [27], [30] and how it could be adapted to fit into the spatial predictive query domain. (ii) Authentication, which introduces techniques to check the completeness and the soundness of the returned answers to users' predictive queries. None of the existing technique [9], [22], [40], [57] can be leveraged directly to handle authentication.

II. Target Audience

The seminar targets researchers, engineers and data scientists in both academia and industry who are working in the fields of mobile data management, spatiotemporal query processing, mobile crowd sourcing, and tracking of moving objects. The seminar summarizes the basic concepts, surveys exiting work, provides industrial application scenarios and envisions the future of spatial predictive queries. More interestingly, the seminar addresses the roadblocks that holds industry back from incubating several techniques that are addressed in the seminar. Attending the seminar does not require prior knowledge about spatial predictive queries. Attendees are expected to gain knowledge about spatial predictive queries both in the Euclidean space and over road networks.

III. BIOGRAPHY

Mohamed Ali is an associate professor at the Institute of Technology, University of Washington, Tacoma. Mohamed's research interests include the processing, analysis and visualization of data streams with geographic and spatial information. For the past decade, Mohamed has been building commercial spatiotemporal data streaming systems to cope with the emerging Big Data requirements. In 2006, Mohamed and his colleagues at the database group at Microsoft Research ramped up the Complex Event Detection and Response (CEDR) project. Then, Mohamed joined the SQL Server group at Microsoft to productize the CEDR project. CEDR has shipped and brand-named as Microsoft StreamInsight. Since the first public release of StreamInsight, Mohamed has been advocating for real-time spatiotemporal data management everywhere; that is the use of StreamInsight in monitoring, managing and mining real time geospatial information across a diversity of verticals. These verticals include but are not limited to: online advertising, behavioral targeting, business intelligence, computational finance, traffic management, social networking, homeland security, emergency and crisis management. In 2011, Mohamed started another journey at Microsoft Bing Maps where he became at the frontline with the Big Data challenge and where he battled various types of spatial search queries. In 2014, Mohamed joined the University of Washington, Tacoma where he leads the geospatial data science team at the Center for Data Science.

Abdeltawab Hendawi is a research associate at the Center for Data Science, University of Washington, Tacoma and a PhD candidate at the Department of Computer Science and Engineering at the University of Minnesota. Abdeltawab's research interests are centred around database systems, big-data mining, spatio-temporal data management, and volunteered geographic information systems. His PhD focuses on predictive query processing against moving objects. Abdeltawab built the iRoad system for predictive queries on road networks, the PANDA system for predictive queries in the Euclidean space, and the iTornado system for predicting the spatio-temporal behavior of severe weather conditions. Prior to joining the University of Minnesota, Abdeltawab obtained his B.Sc. and M.Sc. degrees in Computer Science from Cairo University in Egypt.

REFERENCES

[1] O. Abul, F. Bonchi, and M. Nanni. Never walk alone: Uncertainty for anonymity in moving objects databases. In *Proceedings of the International Conference on Data Engineering, ICDE*, pages 376–385, Cancun, Mexico, Apr. 2008.

[2] S. Ayhan, J. Pesce, P. Comitz, G. Gerberick, and S. Bliesner. Predictive Analytics with Aviation Big Data. In *Proceeding of the ACM SIGSPATIAL International Workshop on Analytics for Big Geospatial Data, BigSpatial*, California, USA, Nov. 2012.

[3] N. Beckmann, H.-P. Kriegel, R. Schneider, and B. Seeger. The R^*-tree: An Efficient and Robust Access Method for Points and Rectangles. In *Proceedings of the ACM International Conference on Management of Data, SIGMOD*, pages 322–331, New Jersey, USA, May 1990.

[4] R. Benetis, C. S. Jensen, G. Karciauskas, and S. Saltenis. Nearest and Reverse Nearest Neighbor Queries for Moving Objects. *VLDB Journal*, 15(3):229–249, 2006.

[5] J. D. L. Blunschi and M. A. V. Salles. Indexing moving objects using short-lived throwaway indexes. *Proceedings of the International Symposium on Advances in Spatial and Temporal Databases, SSTD*, pages 189–207, 2009.

[6] A. Brillingaite and C. S. Jensen. Online Route Prediction for Automotive Applications. In *Proceedings of teh World Congress and Exhibition on Intelligent Transport Systems and Services, ITS*, London, United Kingdom, Oct. 2006.

[7] S. Chen, C. S. Jensen, and D. Lin. A Benchmark for Evaluating Moving Object Indexes. *VLDB Journal*, 1(2):1574–1585, Aug. 2008.

[8] S. Chen, B. C. Ooi, K. L. Tan, and M. A. Nascimento. ST^2B-tree: A Self-tunable Spatio-temporal B^+-tree Index for Moving Objects. In *Proceedings of the ACM International Conference on Management of Data, SIGMOD*, pages 29–42, Vancouver, Canada, June 2008.

[9] W. Cheng and K.-L. Tan. Query assurance verification for outsourced multi-dimensional databases. *Journal of Computer Security*, 17(1):101–126, 2009.

[10] Y.-J. Choi and C.-W. Chung. Selectivity Estimation for Spatio-temporal Queries to Moving Objects. In *Proceedings of the ACM International Conference on Management of Data, SIGMOD*, pages 440–451, Wisconsin, USA, June 2002.

[11] H. D. Chon, D. Agrawal, and A. El-Abbadi. Range and kNN Query Processing for Moving Objects in Grid Model. *Mobile Networks and Applications, MONET*, 8(4):401–412, 2003.

[12] A. Corral, M. Torres, M. Vassilakopoulos, and Y. Manolopoulos. Predictive Join Processing between Regions and Moving Objects. In *Advances in Databases and Information Systems Lecture Notes in Computer Science, ADBIS*, volume 5207, pages 46–61, 2008.

[13] J. Dittrich, L. Blunschi, and M. A. V. Salles. MOVIES: indexing moving objects by shooting index images. *International Journal on Advances of Computer Science for Geographic Information Systems, GeoInformatica*, 15(4):727–767, 2011.

[14] J. Froehlich and J. Krumm. Route Prediction from Trip Observations. In *Society of Automotive Engineers (SAE) World Congress*, Michigan, USA, Apr. 2008.

[15] G. Ghinita, P. Kalnis, A. Khoshgozaran, C. Shahabi, and K.-L. Tan. Private queries in location based services: anonymizers are not necessary. In *Proceedings of the ACM International Conference on Management of Data, SIGMOD*, pages 121–132, Vancouver, Canada, June 2008.

[16] Y. Gu, G. Yu, N. Guo, and Y. Chen. Probabilistic Moving Range Query over RFID Spatio-temporal Data Streams. In *Proceedings of the International Conference on Information and Knowledge Managemen, CIKM*, pages 1413–1416, Hong Kong, China, Nov. 2009.

[17] G. Gulf. Smartphone Users Around the World Statistics and Facts. http://www.go-gulf.com/blog/smartphone/, Jan. 2012.

[18] A. M. Hendawi, J. Bao, and M. F. Mokbel. iRoad: A Framework For Scalable Predictive Query Processing On Road Networks. In *Proceedings of the International Conference on Very Large Data Bases, VLDB*, Riva Del Garda, Italy, Aug. 2013.

[19] A. M. Hendawi, J. Bao, M. F. Mokbel, and M. Ali. Predictive Tree: An Efficient Index for Predictive Queries On Road Networks. In *Proceedings of the International Conference on Data Engineering, ICDE*, Seoul, South Korea, Apr. 2015.

[20] A. M. Hendawi and M. F. Mokbel. Panda: A Predictive Spatio-Temporal

Query Processor. In *Proceedings of the ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, ACM SIGSPATIAL GIS*, California, USA, Nov. 2012.

[21] A. M. Hendawi and M. F. Mokbel. Predictive Spatio-Temporal Queries: A Comprehensive Survey and Future Directions. In *Proceeding of the ACM SIGSPATIAL GIS International Workshop on Mobile Geographic Information Systems, MobiGIS*, California, USA, Nov. 2012.

[22] H. Hu, J. Xu, Q. Chen, and Z. Yang. Authenticating location-based services without compromising location privacy. In *Proceedings of the ACM International Conference on Management of Data, SIGMOD*, pages 301–312, Arizona, USA, May 2012.

[23] H. Hu, J. Xu, and D. L. Lee. A Generic Framework for Monitoring Continuous Spatial Queries over Moving Objects. In *Proceedings of the ACM International Conference on Management of Data, SIGMOD*, pages 479–490, Maryland, USA, June 2005.

[24] C. S. Jensen, D. Lin, and B. C. Ooi. Query and Update Efficient B^{+}-tree Based Indexing of Moving Objects. In *Proceedings of the International Conference on Very Large Data Bases, VLDB*, pages 768–779, Toronto, Canada, Aug. 2004.

[25] H. Jeung, Q. Liu, H. T. Shen, and X. Zhou. A Hybrid Prediction Model for Moving Objects. In *Proceedings of the International Conference on Data Engineering, ICDE*, pages 70–79, Cancn, Mxico, Apr. 2008.

[26] H. Jeung, M. L. Yiu, X. Zhou, and C. S. Jensen. Path Prediction and Predictive Range Querying in Road Network Databases. *VLDB Journal*, 19(4):585–602, Aug. 2010.

[27] P. Kalnis, G. Ghinita, K. Mouratidis, and D. Papadias. Preventing location-based identity inference in anonymous spatial queries. *IEEE Transactions on Knowledge and Data Engineering, TKDE*, 19(12):1719–1733, 2007.

[28] J. Kang, M. F. Mokbel, S. Shekhar, T. Xia, and D. Zhang. Continuous Evaluation of Monochromatic and Bichromatic Reverse Nearest Neighbors. In *Proceedings of the International Conference on Data Engineering, ICDE*, pages 806–815, Istanbul, Turkey, Apr. 2007.

[29] H. A. Karimi and X. Liu. A Predictive Location Model for Location-Based Services. In *Proceedings of the ACM Symposium on Advances in Geographic Information Systems, ACM GIS*, pages 126–133, Louisiana, USA, Nov. 2003.

[30] A. Khoshgozaran and C. Shahabi. Private Information Retrieval Techniques for Enabling Location Privacy in Location-Based Services. In *Proceedings of the International Symposium on Advances in Spatial and Temporal Databases, SSTD*, pages 59–83, Aalborg, Denmark, July 2009.

[31] S.-W. Kim, J.-I. Won, J.-D. Kim, M. Shin, J. Lee, and H. Kim. Path Prediction of Moving Objects on Road Networks Through Analyzing Past Trajectories. In *Proceedings of the International Conference on Knowledge-Based and Intelligent Information and Engineering Systems, KES*, pages 379–389, Vietri sul Mare, Italy, Sept. 2007.

[32] J. Krumm. A markov model for driver turn prediction. *Proceedings of the Society of Automotive Engineers World Congress, SAE*, 2193(1), 2008.

[33] J. Krumm. A survey of computational location privacy. *Personal and Ubiquitous Computing*, 13(6):391–399, 2009.

[34] K. C. K. Lee, H. V. Leong, J. Zhou, and A. Si. An Efficient Algorithm for Predictive Continuous Nearest Neighbor Query Processing and Result Maintenance. In *Proceedings of the International Conference on Mobile Data Management, MDM*, pages 178–182, Ayia Napa, Cyprus, May 2005.

[35] Y. Li, S. George, C. Apfelbeck, A. M. Hendawi, D. Hazel, A. Teredesai, and M. Ali. Routing Service With Real World Severe Weather. In *Proceedings of the ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, ACM SIGSPATIAL GIS*, Texas, USA, Nov. 2014.

[36] M. F. Mokbel, T. M. Ghanem, and W. G. Aref. Spatio-temporal Access Methods. *IEEE Data Engineering Bulletin, DEBull*, 26(2):40–49, June 2003.

[37] M. F. Mokbel, X. Xiong, and W. G. Aref. SINA: Scalable Incremental Processing of Continuous Queries in Spatio-temporal Databases. In *Proceedings of the ACM International Conference on Management of Data, SIGMOD*, pages 443–454, Paris, France, June 2004.

[38] M. F. Mokbel, X. Xiong, M. A. Hammad, and W. G. Aref. Continuous Query Processing of Spatio-temporal Data Streams in PLACE. In *Proceedings of the workshop on Spatio-Temporal Database Management, STDBM*, pages 57–64, Toronto, Canada, Aug. 2004.

[39] L.-V. Nguyen-Dinh, W. G. Aref, and M. F. Mokbel. Spatio-Temporal Access Methods: Part 2 (2003 - 2010). *IEEE Data Engineering Bulletin*, 33(2):46–55, 2010.

[40] S. Papadopoulos, Y. Yang, and D. Papadias. Continuous authentication on relational streams. *VLDB Journal*, 19(2):161–180, 2010.

[41] J. M. Patel, Y. Chen, and V. P. Chakka. STRIPES: An Efficient Index for Predicted Trajectories. In *Proceedings of the ACM International Conference on Management of Data, SIGMOD*, pages 637–646, Paris, France, June 2004.

[42] S. Qiao, C. Tang, H. Jin, T. Long, S. Dai, Y. Ku, and M. Chau. PutMode: prediction of uncertain trajectories in moving objects databases. *Applied Intelligence*, 33(3):370–386, 2010.

[43] K. Raptopoulou, A. Papadopoulos, and Y. Manolopoulos. Fast Nearest-Neighbor Query Processing in Moving-Object Databases. *GeoInformatica*, 7(2):113–137, June 2003.

[44] S. Saltenis and C. S. Jensen. Indexing of Moving Objects for Location-Based Services. In *Proceedings of the International Conference on Data Engineering, ICDE*, pages 463–472, California, USA, Feb. 2002.

[45] S. Saltenis, C. S. Jensen, S. T. Leutenegger, and M. A. Lopez. Indexing the Positions of Continuously Moving Objects. In *Proceedings of the ACM International Conference on Management of Data, SIGMOD*, pages 331–342, Texas, USA, May 2000.

[46] C. Shahabi, L.-A. Tang, and S. Xing. Indexing Land Surface for Efficient kNN Query. In *Proceedings of the International Conference on Very Large Data Bases, VLDB*, pages 1020–1031, Aucklan, New Zealand, Aug. 2008.

[47] A. P. Sistla, O. Wolfson, S. Chamberlain, and S. Dao. Modeling and Querying Moving Objects. In *Proceedings of the International Conference on Data Engineering, ICDE*, pages 422–432, Birmingham U.K, Apr. 1997.

[48] J. Sun, D. Papadias, Y. Tao, and B. Liu. Querying about the Past, the Present, and the Future in Spatio-Temporal. In *Proceedings of the International Conference on Data Engineering, ICDE*, pages 202–213, MASSACHUSETTS, USA, Mar. 2004.

[49] Y. Tao, C. Faloutsos, D. Papadias, and B. L. 0002. Prediction and Indexing of Moving Objects with Unknown Motion Patterns. In *Proceedings of the ACM International Conference on Management of Data, SIGMOD*, pages 611–622, Paris, France, June 2004.

[50] Y. Tao and D. Papadias. Spatial queries in dynamic environments. *ACM Transactions on Database Systems, TODS*, 28(2):101–139, 2003.

[51] Y. Tao, D. Papadias, and J. Sun. The TPR*-tree: An Optimized Spatio-Temporal Access Method for Predictive Queries. In *Proceedings of the International Conference on Very Large Data Bases, VLDB*, pages 790–801, Berlin, Germany, Sept. 2003.

[52] Y. Tao, J. Sun, and D. Papadias. Analysis of predictive spatio-temporal queries. *ACM Transactions on Database Systems, TODS*, 28(4):295–336, Dec. 2003.

[53] Y. Tao, J. Sun, and D. Papadias. Selectivity Estimation for Predictive Spatio-Temporal Queries. In *Proceedings of the International Conference on Data Engineering, ICDE*, pages 417–428, Bangalore, India, Mar. 2003.

[54] H. Tropf and H. Herzog. Multidimensional Range Search in Dynamically Balanced Trees. *Angewante Informatik*, 23(2):71–77, 1981.

[55] H. Wang, R. Zimmermann, and W.-S. Ku. Distributed Continuous Range Query Processing on Moving Objects. In *Proceedings of the International Conference on Database and Expert Systems Applications, DEXA*, pages 655–665, Krakow, Poland, Sept. 2006.

[56] T. Xia and D. Zhang. Continuous Reverse Nearest Neighbor Monitoring. In *Proceedings of the International Conference on Data Engineering, ICDE*, page 77, Georgia, USA, Apr. 2006.

[57] M. L. Yiu, E. Lo, and D. Yung. Authentication of moving kNN queries. In *Proceedings of the International Conference on Data Engineering, ICDE*, pages 565–576, Hannover, Germany, Apr. 2011.

[58] M. L. Yiu, Y. Tao, and N. Mamoulis. The B^{dual}-tree: Indexing Moving Objects by Space Filling Curves in the Dual Space. *VLDB Journal*, 17(3):379–400, May 2008.

[59] M. Zhang, S. Chen, C. S. Jensen, B. C. Ooi, and Z. Zhang. Effectively Indexing Uncertain Moving Objects for Predictive Queries. *PVLDB*, 2(1):1198–1209, 2009.

[60] R. Zhang, H. V. Jagadish, B. T. Dai, and K. Ramamohanarao. Optimized Algorithms for Predictive Range and KNN Queries on Moving Objects. *Information Systems*, 35(8):911–932, Dec. 2010.

[61] R. Zhang, H. V. Jagadish, B. T. Dai, and K. Ramamohanarao. Optimized Algorithms for Predictive Range and KNN Queries on Moving Objects. *Information Systems*, 35(8):911–932, Dec. 2010.