IPHV7I3002X

International Journal Of Advanced Research and Innovation -Vol.7, Issue .III

ISSN Online: 2319 – 9253

Print: 2319 – 9245

ACCESS POLICY CONTROL FOR PROVIDING QOS GUARANTEES IN LARGE NETWORKS

[#1]Badam Veera Prathap - M.Tech Pursuing,
[#2]R.Padma Raj - Assistant Professor,
Department of Computer Science & Engineering,

MOTHER THERESSA COLLEGE OF ENGINEERING & TECHNOLOGY, Peddapalli, Karimnagar, TS, India.

Abstract: - Current event processing systems lack methods to preserve privacy constraints of incoming event streams in a chain of subsequently applied stream operations. This is a problem in large-scale distributed applications like a logistic chain where event processing operators may be spread over multiple security domains. We consider the problem of building online machine-learned models for detecting frauds in e-commence web sites. Since the emergence of the World Wide Web, online shopping and online auction have gained more and more popularity. While people are enjoying the benefits from online trading, criminals are also taking advantages to conduct fraudulent activities against honest parties to obtain illegal profit. Hence proactive fraud-detection moderation systems are commonly applied in practice to detect and prevent such illegal and fraud activities. Machine-learned models, especially those that are learned online, are able to catch frauds more efficiently and quickly than human-tuned rule-based systems. In this paper, we propose an online probit model framework which takes online feature selection, coefficient bounds from human knowledge and multiple instances learning into account simultaneously. By empirical experiments on a real-world online auction fraud detection data we show that this model can potentially detect more frauds and significantly reduce customer complaints compared to several baseline models and the human-tuned rule-based system.

Keywords:- *Data security, Access policy, event processing, security ,multiple instance learning.*

I.INTRODUCTION

In business processes, it is essential to detect inconsistencies or failures early. For example, in manufacturing and logistic processes, items are tracked continuously to detect loss or to reroute them during transport. To answer this need complex event processing (CEP) systems have evolved as a key paradigm for business and industrial applications.CEP systems allow to detect situations by performing operations on event streams which emerge from sensors all over the world, e.g. from packet tracking devices. While, traditionally event processing systems have applied powerful operators in a central way, the emerging increase of event sources and event consumers have raised the need to reduce the communication load by distributed in-network processing of stream operations. In addition, the collaborative nature of today's economy results in large scale networks, where different users, companies, or groups exchange events. As a result, event processing networks are heterogeneous in terms of processing capabilities and technologies, consist of differing participants, and are spread across multiple security domains. However, the increasing interoperability of CEP applications raises the question of security. It is not feasible for a central instance to manage access control for the whole network. Instead, every producer of information should be able to control how its produced data can be accessed. The traditional online shopping business model allows sellers to sell a product or service at a preset price, where buyers can choose to purchase if they find it to be a good deal. Online auction however is a different business model by which items are sold through price bidding. There is often a starting price and expiration time specified by the sellers. Once the auction starts, potential buyers bid against each other, and the winner gets the item with their highest winning bid.

II. RELATED WORK

With the increasing popularity of event-driven systems, a lot of effort has been spent to make the systems secure. For example, a role-based access control is proposed inPesonen et al. and Bacon et al. discuss how publish/subscribe systems can be secured by introducing access control policies in a multi-domain architecture. They describe how event communication between the domains can be supported. Opyrchal et al. present the concept of event owners that can be specified. These are used to provide access to *their* events. Tariq et al. propose a solution to provide authentication and confidentiality in broker-less content-based publish/subscribe systems. Our work is based on the previous work which make event communication secure among different entities in the system. We assume the presence of a system that can handle access control on events. Based on this, we use policy composition in order to derive the necessary access policies at any point during the event processing steps. Access policy composition has found a lot of consideration in distributed systems. Bonatti et al. defined a well recognized algebra for composing access policies. Especially in the area of web service composition, the composition of security policies plays an important role, as different policies have to be combined for every combination of web services (e.g. [23], [24]). We adopt some of these concepts into our distributed CEP system, which allows us to inherit access restrictions during the different processing steps in the operators of our system. To realize our concepts we make use of techniques from statistical inference. More specific, we calculate the Bayesian inference after creating a Bayesian network and learning the dependencies. Since Bayesian inference is a complex calculation, several Monte-Carlo algorithms have been proposed to estimate the inference value(s). They all have in common to arbitrarily pick samples from the Bayesian network probability distribution, and estimate the values based on the samples. The precision of the estimated inference values is dependent on the number of samples. A commonly used technique is the Gibbs sampler.

EXISTING WORK:

On the one hand, sampling techniques can be used to estimate the conditional probabilities of the Bayesian network. However, their precision depends strongly on the number of samples taken from the network, and no approximation scheme exists that allows to draw samples in polynomial time to achieve a certain precision. This makes the approximate algorithms infeasible for security applications, since no guarantees can be made in appropriate time. On the other hand the complexity of calculating *exact* inference can be reduced by storing partial results of the inference calculation which otherwise would have to be calculated multiple times. However, the benefit of these optimizations is heavily dependent on the structure of the Bayesian network.

PROPOSED WORK:

A role-based access control is proposed. Pesonen et al. and Bacon et al. discuss how publish/subscribe systems can be secured by introducing access control policies in a multi-domain architecture. They describe how event communication between the domains can be supported. Opyrchal et al. present the concept of event owners that can be specified. These are used to provide access to *their* events. Tariq et al. propose a solution to provide authentication and confidentiality in broker-less content-based publish/subscribe systems. Our work is based on the previous work which make event communication secure among different entities in the system a. We assume the presence of a system that can handle access control on events. Based on this, we use policy composition in order to derive the necessary access policies at any point during the event processing steps.

III. ACCESS CONTROL FOR CEP

Our approach allows to inherit access requirements by as-signing them to event attributes in form of an *access policy*. This allows to preserve requirements through any chain of dependent correlation steps of operators in *G*. In addition, an obfuscation policy allows to specify an *obfuscation threshold* for event attributes. In each correlation step, the obfuscation of event attributes in produced events is determined by the proposed access policy consolidation protocol. Once the obfuscation threshold is reached for an event attribute, the attribute's access requirements can be ignored. In the following, we detail the concepts behind access policies and obfuscation policies, and formalize the security goal.

A. Access Policies

Access control allows to specify access rights of subjects (operators) for the set of available objects (event attributes). These access rights are provided by the owner of an object (e.g. the producer of an event stream) and may be granted to operators based on an *access requirement*. Such a requirement may be a role, a location or a domain affiliation. Requirements are usually not direct *properties* of the operators, but of the hosts where the operators are deployed. Formally, we specify the access rights within an *access policy AP* for an operator ω as a set of (attribute, access requirement) pairs:

$$AP_{\omega} = \{(att_1, ar_1), ..., (att_n, ar_n)\} .$$

If there is no requirement specified for an attribute, any consumer in the network will be able to access it. Note that we consider attributes to be distinct even if they use the same name, but are produced at two distinct operators.

An access requirement is a tuple of a property *p*, a math-ematical operator *op* and a value set *val*: $ar = (p, op, val)$, where $op \in \{=, <, >, \leq, \geq, \in\}$. *val* can be specified by a range or a set of values. For the sake of simplicity, in this paper access requirements are only referring to domain affiliation and have a structure like this:

$$ar_1 = (domain, \in, \{domainA, domainB\}).$$

In our example scenario, the manufacturer's event attributes have different access requirements. While the information about the item's destination is accessible by the customer, information about where the item is produced and when it can be picked up is restricted to the shipping company. Therefore, the attached AP is defined as follows:

$^{AP}manufacturer$ $^{=}$
{(destination,(domain,∈,{shippingComp,customer})),
(pickup time, (domain,=,shippingComp)),
(production place, (domain,=,shippingComp))}

With the enforcement and assurance of access policies at each producer, a consumer will be eligible to access (receive) an attribute only if the consumer's properties match the access requirements defined for the particular attribute. In this case the consumer is trusted to use the attribute in its correlation function and adopt the requirements specified for the attribute in its own access policy for all produced events.

B. Obfuscation of Event Information

While access policies allow a producer to specify access requirements in a fine-grained manner, the inheritance of requirements in a chain of succeeding operators is at times very restrictive and can limit the efficiency and applicability of the CEP system: in each correlation step of this chain, the number of access requirements may increase by the consolidation of requirements from multiple producers. Each consolidation step can therefore increase the number of interested consumers which are prevented from access to the event attributes of produced event streams. This does not reflect the nature of event processing systems where basic events like single sensor readings may have only little influence on the outcome contained in a complex event representing a specific situation.

In our logistics example, f_{sc} uses *destination*, *production place* and *pickup time* to determine the estimated day of delivery. As a consequence, the customer has no access to the *estimated day of delivery* of the ordered item, since she does not fulfill the access requirements for *production place* and *pickup time*. Yet she has a reasonable interest in this information. And one may claim, that knowledge of the day of delivery does not necessarily allow to draw a relevant conclusion on the *production place* and *pickup time* attribute values. We say, the attribute values get *obfuscated* during the correlation process and depending on the achieved level of obfuscation, the access requirements of an attribute may no longer be needed. In our approach, the level of obfuscation is a measure, to which extent a consumer of the produced attribute (*estimated day of delivery*) can infer the value of the original attribute (*production place*). It can be easily seen in the example, that obfuscation is not only dependent on the values of the attributes, but also on the knowledge of the consumer. Since the *destination* value has led to the *day of delivery* as well, knowledge of the destination would be of great help when trying to infer the restricted attribute *production place* because the delivery time of the item is probably related to the distance between destination and production place. In this work, we will use *obf* $(att_{old}, att_{new}, \omega)$ to refer to the obfuscation achieved by att_{new} for att_{old} given the knowledge available at a consumer $\omega \in \Omega$.

We allow every operator to specify with its access policy also an obfuscation policy. The obfuscation policy contains obfuscation thresholds for the attributes the operator produces. During the processing of an event attribute, its obfuscation w.r.t. each potential consumer is calculated. Once, the obfuscation threshold for a consumer is reached, the event attribute can be delivered in spite of conflicting access requirements. Formally, we define the obfuscation policy *OP* for an operator ω as a set of (attribute, obfuscation threshold) pairs:

$$OP_\omega = \{(att_1, ot_1), ..(att_n, ot_n)\} .$$

For instance, the obfuscation policy

$$OP_{manufacturer} = \{(destination, 0.9)\}.$$

allows the shipping company for events addressed to the consumer to ignore all access rights for *destination* in the access policy of attribute *day of delivery* if *obf* (*destination, day of delivery,* ω_C) ≥ 0.9. We detail the exact semantics of the obfuscation value and its measure in Section IV.

C. Security Goal

Let $att_{old} \rightarrow_\omega att_{new}$ denote that

1) at some operator $\omega \in \Omega$, att_{old} is taken as input to the correlation function f_ω , and

2) f_ω produces att_{new} in dependence of att_{old}.

Furthermore, let $att_{old} \rightarrow^* att_{new}$ denote the transitive closure of the dependency relation. For any pair of attributes with $att_{old} \rightarrow^* att_{new}$ we say that att_{new} is *dependent* on att_{old}. Our main goal is to preserve the privacy of event attributes over multiple correlation steps by respecting the dependency relationship between the attributes produced by the CEP system. In particular, access requirements must not be applied solely to the attribute att_{old}, but have to be inherited to all dependent attributes (att_{new}) unless a sufficient obfuscation threshold for att_{new} has been reached.

More formally, given for each attribute *att* an initial set of access requirements denoted by $AR_{init}(att)$. We require for any policy consolidation algorithm two conditions to be met:

Condition 1. *For all attributes att* $\in O_\omega$ *produced at* ω

$$AR_{init}(att) \subset AP_\omega . \qquad (1)$$

Condition2. *For all dependent attribute pairs* $(att_{old}, att_{new}) \in \rightarrow^*$ *with*

1) ω_i *has produced* att_{old} *with access requirement* $AR(att_{old})$ *and obfuscation threshold* $(att_{old}, x) \in$ OP_{ω_i} ,

2) att_{new} *is produced by* ω_j

3) att_{new} *is consumed by* ω_k

the access requirement in $AP_{\omega j}$ *yield*

$$AR(att_{old}) \subset AP_{\omega j} \text{ if } obf\,(att_{old}, att_{new}, \omega_k) < x. \qquad (2)$$

A policy consolidation algorithm needs to ensure Condi-tion 1 and Condition 2 in the presence of adversaries who try to derive event attribute values they are by policy not allowed to access directly. We want to avoid that hosts maliciously or inadvertently obtain information from event streams for which they have no authorization. Note, by accessing event streams according to the specified system model, hosts may still be able to infer event attributes of unauthorized event streams from legally received event streams. An adversary in our system is therefore limited to the behavior described in the system model. The adversary is authenticated and can only access streams according to its properties. The derived event output follows the operator specification and the access require-ments for each executed operator. Each adversary is bound to analyzing outgoing event streams which it is allowed to access, for inferring any additional information.

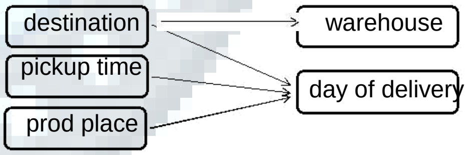

Figure1. Dependency Graph of the Shipping Company Operator

IV. POLICY CONSOLIDATION AND EVENT OBFUSCATION

To meet the security goal from Section III our approach establishes secure event streams between each pair of operators in G. For establishing secure event streams we rely on mechanisms available in state of the art publish/subscribe systems including our own work.our approach it is only important to understand that each consumer ω_c needs to request required event attributes. The requests are handled at the producer ω_p and ω_c will need to authenticate itself against ω_p for the corresponding event attribute. After successful authentication ω_p will forward to ω_c

1) only those events matching the request of ω_c,
2) only those events containing attributes att s.t.
   a) the access policy of att allows ω_c access to att,
   b) att has achieved a sufficiently high obfuscation,

i.e. $\forall (att_i, ot_i) \in OP_{\omega p}\ obf(att_i, att, \omega_c) \geqslant ot_i$

To this end ω_p will have to perform on its incoming streams an access policy consolidation to ensure all necessary access policies can be inherited and a calculation of the obfuscation values $obf(att_i, att, \omega_c)$. In the following we will show the approach to access consolidation by modeling all potential dependencies between incoming and outgoing event streams in an event dependency graph and calculate obfuscation policies by relying on a Bayesian network.

V. Scalable Access Policy Consolidation

Instead of accounting for a global Bayesian network, we propose to exploit local knowledge available at each host. This allows us to reduce the number of relations of incoming and outgoing attributes and thus leads to a huge gain in processing overhead. The idea of our approach is that a host in the CEP network creates a local Bayesian network for each of its deployed operators. The handling (i.e. forwarding) of the event is based on the locally achieved obfuscation. This limits the computational effort by accepting that obfuscation is not measured over multiple correlation steps, and therefore some events may be treated more restrictive than actually needed.

Algorithm 1 Local Obfuscation Calculation

```
procedure INITIALIZE($\omega$) for all operator $\omega$ do
    $D_\omega \leftarrow$ FINDMULTIPATHOPERATORS($\omega$) end for
    for all $\omega \in D_\omega$ do
        $relAtts \leftarrow$ FINDRELATEDATTRIBUTES for all $(att_{new}, att_{old}) \in relAtts$ do
            TRANSMIT P($att_{new} / att_{old}$) TO $\omega$ end for
    end for end procedure
procedure UPONRECEIVEEVENT($e$) for all $att \in e$ do
    if $\exists$ multPathDependency(att) then
        CALCULATEWORSTCASEOBFUSCATION(ATT)
    else
        CALCULATELOCALOBFUSCATION(ATT) end if
    end for end procedure
```

A. Measuring Local Obfuscation

In the approach, every host calculates obfuscation only for the locally known attribute dependencies (i.e. $att_{old} \rightarrow_{\omega} att_{new}$) in contrast to calculating the obfuscation for every pair of dependent attributes (i.e. $att_{old} \rightarrow^{*} att_{new}$). This has three major benefits: i) a smaller dependency graph, ii) less communication overhead, and iii) the network is not multiply connected, because there exist only paths of length 1. As a consequence, every host can create a *local dependency graph* on its own instead of creating a global dependency graph for all dependent attributes. Furthermore, we can efficiently calculate the exact inference probabilities by applying variable elimination optimization for single connected networks to efficiently determine the obfuscation value (cf. Section IV-B).

Even in a local approach for obfuscation calculation the multi-path dependencies of attributes need to be considered. Attributes might reach the recipient via multiple paths (i.e. parallel chains of operators in a multiply-connected correlation network, cf. Figure 5). An adversary that can subscribe to such attributes may be able to infer the original value by combining the event information received through the multiple paths. We meet this by analyzing the entire operator graph during initialization of our algorithm (c.f. Algorithm 1). For every attribute pair with multi-path dependencies the operators that reside on distinct paths exchange the dependency functions w.r.t. the attributes. For example, in a scenario as depicted in Figure 5, the inference probability is calculated as follows:

$$P\,(att_{old}|att_1, att_2) = \alpha * P(att_{old}) * P(att_1|att_{old}) * P(att_2|att_{old}) \qquad (5)$$

where α is the normalization constant $1/P(att_2)$.

Hence, $P(att_1|att_{old})$ is sent to operator ω_2 and $P(att_2|att_{old})$ to operator ω_1 vice versa.

After performing the initialization, each operator can calculate the obfuscation value from local knowledge only. In the above example, if operator ω_1 now calculates the obfuscation of an incoming attribute att_{old} for the outgoing attribute att_1, it uses the dependency functions received during the initialization phase. There, it searches for the outcome att_2 which has the highest chance for inferring att_{old}, i.e. the entry with the highest probability. This value is then used in the calculation of $P(att_{old}|att_1, att_2)$, as it results in the minimal achievable obfuscation.

Note that the initialization needs to be performed with each change of the correlation graph and follows the learning phase of the Bayesian networks. However, changes to the operator graph typically are for many practical settings a result of changes to the business logic. Hence, we expect only rare interruptions of the event processing service.

B. Correctness

As our work addresses mainly how to establish producer centric access policies in CEP in a scalable way, we give only informal correctness arguments under the limitations for the adversary introduced in Section III. Three main properties guarantee that the proposed approach is correct in terms of the defined security goal:

1) According to our assumptions in Section III, an adver-sary tries to infer additional information by analyzing all event streams which it is allowed to access. The proposed algorithm considers the complete knowledge the consumer *might* have. That means, it is considered that every attribute influencing the requested local obfuscation ($obf(att_{old}, att_{new}, \omega_c)$) that is accessible to the consumer is known.
2) In accordance to Property 1, every path from att_{old} to att_{new} is considered in the algorithm. That means, every piece of information an adversary may access in order to infer att_{old} is included when calculating the inference.
3) Locally unknown events (which may occur in multi-path dependency calculations) are always handled as a worst-case-consideration. We always use the value in our calculations which would give an adversary the most inference information, i.e. the value resulting in the worst obfuscation.

Since all sources of event information which might influence the obfuscation value of any operator are considered in our approach, the obfuscation value calculated at an operator cannot further be lowered by any means. Hence, with the presented approach, we guarantee: If the consumer does not fulfill the access requirements for an attribute att_{old}, it will also not be able to access any attribute att_{new} if the attributes depend on each other ($(att_{old} \xrightarrow{*} att_{new})$ unless a sufficient obfuscation threshold for att_{new} has been reached. We do not guarantee, though, that the consumer will receive every attribute that has achieved a sufficient obfuscation.

V.PROPOSED TECHNIQUES

• **Rule-based features:**

Human experts with years of experience created many rules to detect whether a user is fraud or not. An example of such rules is “blacklist”, i.e. whether the user has been detected or complained as fraud before. Each rule can be regarded as a binary feature that indicates the fraud likeliness.

• **Selective labeling:**

If the fraud score is above a certain threshold, the case will enter a queue for further investigation by human experts. Once it is reviewed, the final result will be labeled as boolean, i.e. fraud or clean. Cases with higher scores have higher priorities in the queue to be reviewed. The cases whose fraud score are below the threshold are determined as clean by the system without any human judgment.

• **Fraud churn:**

Once one case is labeled as fraud by human experts, it is very likely that the seller is not trustable and may be also selling other frauds; hence all the items submitted by the same seller are labeled as fraud too. The fraudulent seller along with his/her cases will be removed from the website immediately once detected.

VI.CONCLUSION

This paper addressed the inheritance and consolidation of access policies in heterogeneous CEP systems. We identified a lack of security in multi-hop event processing networks and proposed a solution to close this gap. More specific, we presented an approach that allows the inheritance of access requirements, when events are correlated to complex events. Our algorithm includes the obfuscation of information, which can happen during the correlation process, and uses the obfuscation value as a decision-making basis whether inheritance is needed. We presented an implementation of our approach, based on Bayesian Network calculations.

In this paper we build online models for the auction fraud moderation and detection system designed for a major Asian online auction website. By empirical experiments on a realword online auction fraud detection data, we show that our proposed online probit model framework, which combines online feature selection, bounding coefficients from expert knowledge and multiple instance learning, can significantly improve over baselines and the human-tuned model. Note that this online modeling framework can be easily extended to many other applications, such as web spam detection, content optimization and so forth. Regarding to future work, one direction is to include the adjustment of the selection bias in the online model training process. It has been proven to be very effective for offline models in. The main idea there is to assume all the unlabeled samples have response equal to 0 with a very small weight. Since the unlabeled samples are obtained from an effective operation system, it is reasonable to assume that with high robabilities they are non-fraud.

REFERENCES

[1] A. Buchmann and B. Koldehofe, “Complex event processing,” *it - Information Technology*, vol. 51:5, pp. 241–242, 2009.

[2] Access Policy Consolidation for Event Processing Systems Bj¨orn Schilling∗, Boris Koldehofe∗, Kurt Rothermel∗ and Umakishore Ramachandran†∗*Institute for Parallel and Distributed Systems, Universit¨at Stuttgart.*

[3] A. Hinze, K. Sachs, and A. Buchmann, “Event-based applications and enabling technologies,” in *Proceedings of the Third ACM International Conference on Distributed Event-Based Systems*, ser. DEBS ’09. New York, NY, USA: ACM, 2009, pp. 1:1–1:15.

[4] P. Pietzuch, “Hermes: A scalable event-based middleware,” Ph.D. dissertation, University of Cambridge, 2004.

[5] G. Li and H.-A. Jacobsen, “Composite subscriptions in content-based publish/subscribe systems,” in *Proc of*

the 6th Int. Middleware Conf., 2005, pp. 249–269.

[6] G. G. Koch, B. Koldehofe, and K. Rothermel, "Cordies: expressive event correlation in distributed systems," in *Proc. of the 4th ACM International Conference on Distributed Event-Based Systems (DEBS)*, 2010, pp. 26–37.

[7] B. Koldehofe, B. Ottenwalder,¨ K. Rothermel, and U. Ra-machandran, "Moving range queries in distributed complex event processing," in *Proc. of the 6th ACM International Conference on Distributed Event-Based Systems (DEBS)*, 2012, pp. 201–212.

[8] B. Schilling, B. Koldehofe, U. Pletat, and K. Rothermel, "Distributed heterogeneous event processing: Enhancing scalability and interoperability of CEP in an industrial context," in *Proc. of the 4th ACM International Conference on Distributed Event-Based Systems (DEBS)*, 2010, pp. 150–159.

[9] B. Schilling, B. Koldehofe, and K. Rothermel, "Efficient and distributed rule placement in heavy constraint-driven event systems," in *Proc. of the 10th IEEE International Confer-ence on High Performance Computing and Communications (HPCC)*, 2011, pp. 355–364.

[10] M. A. Tariq, B. Koldehofe, A. Altaweel, and K. Rothermel, "Providing basic security mechanisms in broker-less pub-lish/subscribe systems," in *Proceedings of the 4th ACM Int. Conf. on Distributed Event-Based Systems (DEBS)*, 2010, pp. 38–49.

[11]. L. I. W. Pesonen, D. M. Eyers, and J. Bacon, "Encryption-enforced access control in dynamic multi-domain publish/subscribe networks," in *Proc. of the 2007 ACM International Conference on Distributed Event-Based Systems (DEBS)*, 2007, pp. 104–115.

[12]J. Bacon, D. M. Eyers, J. Singh, and P. R. Pietzuch, "Access control in publish/subscribe systems," in *Proc. of the 2nd ACM International Conference on Distributed Event-Based Systems (DEBS)*, 2008, pp. 23–34.

[13]M. A. Tariq, B. Koldehofe, G. G. Koch, I. Khan, and K. Rothermel, "Meeting subscriber-defined QoS constraints in publish/subscribe systems," *Concurrency and Computation: Practice and Experience*, vol. 23, no. 17, pp. 2140–2153, 2011.

[14]S. Rizou, F. Durr,¨ and K. Rothermel, "Providing qos guaran-tees in large-scale operator networks," in *High Performance Computing and Communications (HPCC), 2010 12th IEEE International Conference on*, 2010, pp. 337 –345.

[15]S. J. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach, 2nd ed.* Prentice Hall, 2002.

[15]S. Geman and D. Geman, "Stochastic relaxation, gibbs dis-tributions, and the bayesian restoration of images," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. PAMI-6, pp. 721 –741, 1984.

[16].A. E. Gelfand and A. F. M. Smith, "Sampling-based approaches to calculating marginal densities," *Journal of the American Statistical Association*, vol. 85, no. 410, pp. 398–409, 1990.