

Upgrading from Gaussian Processes to Student's-T Processes

Brendan D. Tracey *

*Santa Fe Institute, Santa Fe, NM, 87501
Massachusetts Institute of Technology, Cambridge, MA, 02139*

David H. Wolpert †

*Santa Fe Institute, Santa Fe, NM, 87501
Massachusetts Institute of Technology, Cambridge, MA, 02139
Arizona State University, Tempe, AZ, 85287*

Gaussian processes are the most commonly deployed tool in aerospace design for performing Bayesian optimization. Despite their success, Gaussian processes suffer significant drawbacks. Outliers are assumed to be unlikely, due to the low kurtosis of the Gaussian distribution, and the posterior variance depends only on the observed locations, not the observed function values. Student's-T processes have seen recent development as a generalization of Gaussian processes. They have the more general Student's-T distribution as a foundation instead of the Gaussian distribution. Student's-T processes maintain the primary advantages of Gaussian processes (kernel function, analytic update rule), and are additionally much more likely to generate outliers, and the posterior uncertainty increases or decreases depending on the variance of the observed values. This paper describes the Student's-T process in the context of aerospace optimization. It shows how to construct a Student's-T process using a kernel function, to update the process given new samples, and to compute optimization-relevant quantities such as expected improvement. Finally, it compares the performance of the Student's-T process with a Gaussian process on canonical test problems in Bayesian optimization, and applies the Student's-T process to the optimization of an aerostructural design problem.

I. Nomenclature

D	=	Data set of inputs and observed function values at those inputs
$\mathbb{E}[x]$	=	Expected value of random variable x
EI	=	Expected Improvement
f	=	Objective function of interest
GP	=	Gaussian process
K	=	Matrix of pairwise kernel function evaluations
$k(x, x')$	=	Kernel function of a process
$m(x)$	=	Mean function of a process
MVG	=	Multivariate Gaussian distribution
MVT	=	Multivariate Student's-T distribution
$p(y x)$	=	Marginal probability of y at an input x
$p(y x, D)$	=	Marginal probability of y at an input x given data D
STP	=	Student's-T process
\mathcal{T}	=	Multivariate Student's-T distribution
x	=	Input to the objective function f
\tilde{x}	=	Observed input to the objective function f (with corresponding \tilde{y})
y	=	Output from the objective y
y^*	=	Global optimum of a function
\hat{y}	=	Best current known function value
\tilde{y}	=	Observed output from the objective y
μ	=	Mean parameter of the distribution

*Postdoctoral Researcher, tracey.brendan@gmail.com

†Resident Faculty, dhw@santafe.edu

ν	=	Degrees of freedom parameter for the Student's-T distribution/process
Σ	=	Shape parameter of the distribution
σ	=	Marginal shape parameter at a given input
ξ	=	Constraint violation

II. Introduction

One major challenge in aerospace design optimization is the computational expense in determining the quantity of interest. An Euler flow solution around a full aircraft geometry is still a non-trivial endeavor, let alone computing the aerodynamic forces with higher-fidelity governing equations like RANS and DES. These higher fidelity simulations are being increasingly relied upon by the design community, a trend which will only increase as researchers work to reduce the current deficiencies in computational fluid dynamics (CFD), such as turbulent transition, separation, and engine stall analysis [1]. These acute computational requirements place a premium on data efficiency when designing optimization algorithms.

Bayesian optimization is a class of global optimization techniques designed for expensive objective functions. Their utility is in the efficiency with which they discover global minima of a function, in contrast to local hill-descending methods such as the Nelder-Mead simplex [2] or quasi-Newton methods [3] that only find local minima, and in contrast to population- or sampling-based global methods such as genetic algorithms [4] or simulated annealing [5] that typically require hundreds or thousands of simulations. Bayesian optimization techniques work by defining a relative likelihood of different function behaviors in the form of a probabilistic *prior* over the space of functions. As the objective function is evaluated, this prior distribution can be updated using Bayes rule to get an updated belief, i.e. *posterior*, about the function behavior at unobserved inputs. This Bayesian update is often computationally expensive relative to optimization algorithms, but this added expense is often negligible relative to evaluating a flow solution.

The most common prior used in Bayesian optimization is that of a Gaussian process (GP). A Gaussian process assumes that the marginal uncertainty over any (finite) set of input locations is a multivariate Gaussian distribution (MVG). Gaussian processes have several desirable mathematical properties, importantly the Bayesian update step is analytic given new objective function evaluations, as is finding the marginal distributions for the function behavior at unknown locations. Gaussian processes have been widely used in global optimization, from the important EGO algorithm [6], to supersonic jet optimization [7], to multi-fidelity modeling and optimization [8, 9]. In spite of their success, GPs have several known shortcomings. First, the Gaussian distribution is not a “heavy-tailed” distribution, and so extreme outliers are unlikely to occur. Second, the posterior variance of a Gaussian process does not depend on the returned objective values, but depends *only* on the input evaluation locations. This implies, for instance, that the posterior uncertainty is no higher if the returned objective values were all much different than expected.

This paper argues for the use of a different probabilistic prior — a Student's-T process (STP). Student's-T processes assume that any (finite) set of input locations are jointly distributed according to a multivariate Student's-T distribution (MVT), unlike the MVG of a Gaussian process. Like a GP, there is an analytic formula for updating an STP with new samples, and it is easy to find the marginal distribution for unknown locations. Additionally, the Student's-T distribution includes an additional “degrees of freedom” parameter controlling the kurtosis of the distribution. This means outliers can be much more likely under an STP than a GP, and additionally, as will be shown, the posterior variance is increased if observed values are different than expected. As the degrees of freedom parameter approaches infinity the STP converges to a GP, and so STPs are a generalization of GPs.

In this manuscript, we begin by reviewing Gaussian processes. We then introduce the Student's-T process, comparing and contrasting it with GPs. We then derive formulae useful for Bayesian optimization, such as expected improvement and marginal likelihood. Finally, we compare the performance of GPs with STPs on benchmark optimization problems and an aerostuctural design optimization problem. We find that the Student's-T processes significantly outperform their Gaussian counterparts.

III. Gaussian Process Review

In this section we briefly review Gaussian processes not only as an aid to the reader, but also to ease the comparison with Student's-T processes described in the next section.

As described in the introduction, Gaussian processes place a probabilistic prior over the space of functions. They are defined as a “collection of random variables, any finite number of which have a joint Gaussian distribution”[10]. GPs are parameterized by two functions. The first is a mean function, $m(x)$, that sets the prior expected value of any

input location x . The second is a kernel function $k(x, x')$, that sets the covariance between two different function inputs x and x' . Together, these functions provide a probabilistic description of the objective function:

$$f(x) \sim \mathcal{GP}(m(x), k(x, x')) \quad (1)$$

At a single input location, x , the prior uncertainty of $f(x)$ is described by a (univariate) Gaussian distribution

$$p(y|x) = \mathcal{N}(\mu, \sigma) \quad (2)$$

where $\mu = m(x)$, and $\sigma = k(x, x)$. At multiple input locations, their joint marginal uncertainty is described by a multivariate Gaussian

$$p(y|x) = \mathcal{N}(\mu, \Sigma) \quad (3)$$

where $\mu_i = m(x_i)$, and $\Sigma_{i,j} = k(x_i, x_j)$. It is often the case that $m(x)$ is assumed to be 0, and this will be assumed for the remainder of the work. The kernel function is typically chosen so that the covariance decreases as x and x' become farther apart. A common choice is the ‘‘isotropic squared exponential’’ function (SE)

$$k(x, x') = \exp(-0.5\|x - x'\|_2^2/\sigma^2) \quad (4)$$

where σ is a ‘‘bandwidth’’ hyperparameter that sets the correlation length.

This prior can be updated when the objective function is evaluated at a given x , and the output value $y = f(x)$ is observed. An important property of the Gaussian distribution is that it is closed under conditioning, i.e. the posterior distribution after observing a set of samples is still a Gaussian distribution in the other variables. It can be shown that given a set of inputs and outputs $D = \{\{\tilde{x}_1, \tilde{y}_1\}, \{\tilde{x}_2, \tilde{y}_2\}, \{\tilde{x}_3, \tilde{y}_3\}, \dots\}$, the uncertainty for a set of unobserved locations x is given by

$$p(y|x, D) = \mathcal{N}(\mu, \Sigma) \quad (5)$$

$$\mu = K_{x, \tilde{x}} K_{\tilde{x}, \tilde{x}}^{-1} \tilde{y} \quad (6)$$

$$\Sigma = K_{x, x} - K_{x, \tilde{x}} K_{\tilde{x}, \tilde{x}}^{-1} K_{\tilde{x}, x} \quad (7)$$

where $K_{\tilde{x}, \tilde{x}}$ is the covariance matrix defined by the kernel function between the observed locations in D , $K_{\tilde{x}, x}$ is the covariance between the observed locations and the unobserved locations, and $K_{x, x}$ is the covariance among the unobserved locations. It is notable that the posterior covariance matrix (7) *only* depends on the observed locations, \tilde{x} , and not on the function values themselves \tilde{y} .

GPs are often used for function minimization, where the goal is to find an input location with a low function response. One common strategy is to choose the next input to query by selecting the input with the highest *expected improvement* beyond the current best known value. Given the current best function value, \hat{y} , the next x chosen as

$$\operatorname{argmax}_x EI(x) = \operatorname{argmax}_x \int 1\{y < \hat{y}\}(\hat{y} - y)p(y|x, D)dy \quad (8)$$

where $p(y|x, D)$ is computed according to (5). The expected improvement under Gaussian uncertainty has an analytic form

$$EI(x) = (\hat{y} - \mu)\Phi(z) + \sigma * \phi(z) \quad (9)$$

where $z = (\hat{y} - \mu)/\sigma$, and Φ and ϕ are the cumulative density function (CDF) and probability density (PDF) for the Gaussian distribution respectively. This analytic expression for expected improvement makes it relatively efficient to solve the optimization problem in (8).

IV. Student’s-T Processes

Student’s-T processes are an alternate prior for Bayesian optimization. Just as GPs have marginal distributions described by the multivariate Gaussian distribution, STPs have marginal distributions described by the multivariate Student’s-T distribution [11]. Student’s-T processes receive mention in [10] and have been used occasionally in modeling [12, 13]. However, a renewal of interest began with Shah et. al. [14] in which the authors derive the Student’s-T process

from a Wishart prior, and show that STPs are the most general elliptic process with analytic density. The process has been further explored for Bayesian optimization [15].

The multivariate Student's-T distribution is a generalization of the multivariate Gaussian distribution with an additional parameter, ν , describing the degrees of freedom of the distribution. The probability density is given by:

$$\mathcal{T}(\mu, \Sigma, \nu) = \frac{\Gamma((\nu + d)/2)}{\Gamma(\nu/2)\nu^{d/2}\pi^{d/2}|\Sigma|^{1/2}} \left(1 + \frac{1}{\nu}(y - \mu)^T \Sigma^{-1}(y - \mu)\right)^{-(\nu+d)/2} \quad (10)$$

where d is the dimension of the distribution, μ is a location parameter, Σ is a shape parameter, and $\nu > 2$ is the degrees of freedom. Like in a Gaussian distribution, μ is the mean (and mode) of the distribution. The shape parameter Σ is not the covariance matrix of the distribution, but is related to it,

$$E[(y - \mu)^T (y - \mu)] = \frac{\nu}{\nu - 2} \Sigma \quad (11)$$

As the degrees of freedom increases, i.e. $\nu \rightarrow \infty$, the multivariate Student's-T distribution converges to a multivariate Gaussian distribution with the same mean and shape parameter.

The Student's-T process, as before, is defined with a mean function and a kernel function, and additionally defined by the degrees of freedom

$$f(x) \sim \mathcal{STP}(m(x), k(x, x'), \nu) \quad (12)$$

The mean function $m(x)$ defines the prior expected value at each location, x . The kernel function, as in a GP, sets the *covariance* between two locations x and x' . Thus, the joint distribution for a finite subset of locations is

$$p(y|x) = \mathcal{T}(\mu, \Sigma, \nu) = \mathcal{T}\left(\mu, \frac{\nu - 2}{\nu} K, \nu\right) \quad (13)$$

where μ is the vector of means, $\mu_i = m(x_i)$, and K is the matrix of pairwise-kernel evaluations $K_{i,j} = k(x_i, x_j)$ (remembering the covariance is not the same as the shape parameter).

The Student's-T process has an analytic Bayesian update rule as new data is gathered. The multivariate Student's-T distribution, and by extension the Student's-T process, is closed under conditioning. It can be shown [14, 16] that given a set of samples, $D = \{\{\tilde{x}_1, \tilde{y}_1\}, \{\tilde{x}_2, \tilde{y}_2\}, \{\tilde{x}_3, \tilde{y}_3\}, \dots\}$, the posterior distribution is given by

$$p(y|x, D) = \mathcal{T}\left(\hat{\mu}, \frac{\nu - 2}{\nu} \hat{K}, \hat{\nu}\right) \quad (14)$$

$$\hat{\mu} = K_{x,x} K_{\tilde{x},\tilde{x}}^{-1} \tilde{y} \quad (15)$$

$$\hat{K} = \frac{\nu + \tilde{y}^T K_{\tilde{x},\tilde{x}}^{-1} \tilde{y} - 2}{\nu + |D| - 2} \left(K_{x,x} - K_{x,\tilde{x}} K_{\tilde{x},\tilde{x}}^{-1} K_{\tilde{x},x}\right) \quad (16)$$

$$\hat{\nu} = \nu + |D| \quad (17)$$

The posterior mean for an STP is identical to that of a GP (for the same kernel function), seen by comparing (6) and (15). The posterior covariance, however, is different between the two processes. The rightmost term in (16) matches its Gaussian counterpart, but the leading term has no equivalent in a GP. We see that this term has an explicit dependence on \tilde{y} , the function outputs, and scales the posterior covariance. The important relation is

$$\tilde{y}^T K_{\tilde{x},\tilde{x}}^{-1} \tilde{y} \stackrel{?}{=} |D| \quad (18)$$

which compares the squared Mahalanobis distance of the samples using their covariance, and compares it with the number of observed samples. If $\tilde{y}^T K_{\tilde{x},\tilde{x}}^{-1} \tilde{y} > |D|$, then the posterior covariance is larger than its Gaussian counterpart, while if less, the posterior covariance is correspondingly lower. These two terms are equal in expectation *if* the output values are actually generated from a Gaussian processes, i.e. $y \sim \mathcal{N}(0, \Sigma)$. The squared Mahalanobis distance for Gaussian-generated samples is distributed according to a χ^2 distribution with a mean $|D|$ [17]. This implies that if the observed y values vary from each other by about as much as one would expect under a GP, then the posterior covariance under the STP is roughly identical to the covariance of the equivalent GP. On the other hand, if the values vary by significantly more or less than expected, then the STP adapts accordingly. Note, however, that as, $\nu \rightarrow \infty$, this extra term tends to unity. As expected, the difference in the STP prediction is most prominent for small values of ν .

The benefits of using an STP do not come with a significantly higher computational cost. The only major computational difference between the two update rules is the additional term in (16). However, for both processes, the dominant cost is computing the Cholesky decomposition of $\Sigma_{\tilde{x},\tilde{x}}$ to compute terms involving $K_{\tilde{x},\tilde{x}}^{-1}$, which can be computed once and cached. In fact, the expression $K_{\tilde{x},\tilde{x}}^{-1}\tilde{y}$ occurs in both (15) and (16), so if the mean prediction has already been computed, computing the additional scaling factor of the STP only requires computing a dot product, which is $O(|D|)$.

A. Expected Improvement

The marginal distribution for the output at a single (unobserved) input is a Student's-T distribution, as discussed above. The Student's-T distribution also has an analytic expression for the expected improvement over a given \hat{y} , which we derive here. This implies that selecting the next design location using expected improvement is equivalently easy under an STP to under a GP.

The probability density for a Student's-T distribution is given by

$$\mathcal{T}(\mu, \sigma, \nu) = \frac{\Gamma(\frac{\nu+1}{2})}{\sqrt{\nu\pi}\Gamma(\frac{\nu}{2})} \frac{1}{\sigma} \left(1 + \frac{((y-\mu)/\sigma)^2}{\nu}\right)^{-\frac{\nu+1}{2}} \quad (19)$$

For simplicity, let

$$C = \frac{\Gamma(\frac{\nu+1}{2})}{\sqrt{\nu\pi}\Gamma(\frac{\nu}{2})} \quad (20)$$

The expected improvement equation becomes

$$\int_{-\infty}^{\infty} 1\{y < \hat{y}\}(\hat{y} - y)p(y|x)dy = \int_{-\infty}^{\hat{y}} (\hat{y} - y) \frac{1}{\sigma} C \left(1 + \frac{((y-\mu)/\sigma)^2}{\nu}\right)^{-\frac{\nu+1}{2}} dy \quad (21)$$

First, substitute $t = \frac{y-\mu}{\sigma}$, $dt = dy/\sigma$.

$$= \int_{-\infty}^{\frac{\hat{y}-\mu}{\sigma}} (\hat{y} - \mu - \sigma t) C \left(1 + \frac{t^2}{\nu}\right)^{-\frac{\nu+1}{2}} dt \quad (22)$$

$$= (\hat{y} - \mu) \int_{-\infty}^{\frac{\hat{y}-\mu}{\sigma}} C \left(1 + \frac{t^2}{\nu}\right)^{-\frac{\nu+1}{2}} dt - \int_{-\infty}^{\frac{\hat{y}-\mu}{\sigma}} \sigma t C \left(1 + \frac{t^2}{\nu}\right)^{-\frac{\nu+1}{2}} dt \quad (23)$$

Note that the inside of the left integral is the formula for the standard Student's-T distribution (a Student's-T distribution with $\mu = 0$ and $\sigma = 1$), so this integral is just the CDF of the standard T distribution. For convenience, make the substitution $z = (\hat{y} - \mu)/\sigma$. Simultaneously, in the right integral make the substitution $s = 1 + t^2/\nu$, so $t dt = \nu/2 ds$.

$$= (\hat{y} - \mu)\Phi_s(z) - \sigma \int_{-\infty}^{1+\frac{z^2}{\nu}} \frac{\nu}{2} C s^{-\frac{\nu+1}{2}} ds \quad (24)$$

$$= (\hat{y} - \mu)\Phi_s(z) - \sigma \frac{\nu}{2} \frac{2}{1-\nu} C s^{-\frac{\nu+1}{2}} \Big|_{-\infty}^{1+\frac{z^2}{\nu}} \quad (25)$$

$$= (\hat{y} - \mu)\Phi_s(z) + \sigma \frac{\nu}{\nu-1} \left(1 + \frac{z^2}{\nu}\right) C \left(1 + \frac{z^2}{\nu}\right)^{-\frac{\nu+1}{2}} \quad (26)$$

$$= (\hat{y} - \mu)\Phi_s(z) + \frac{\nu}{\nu-1} \left(1 + \frac{z^2}{\nu}\right) \sigma \phi_s(z) \quad (27)$$

The formula (27) is the expected improvement over \hat{y} for a Student's-T distributed random variable, where again $z = (\hat{y} - \mu)/\sigma$ and Φ_s and ϕ_s are the CDF and PDF of the standard Student's-T distribution respectively. As expected, this converges to (9) as $\nu \rightarrow \infty$. When ν is small, however, the expected improvement can be significantly different from a Gaussian, even when both distributions have the same mean and standard deviation. A smaller ν increases the likelihood of outliers, and so more importance is placed on having a large uncertainty than on having a promising mean.




Fig. 1 Comparison of (a) a Gaussian process and (b) a Student's-T process with the same prior mean and kernel function. Notice that extreme outliers are much more likely under the Student's-T process than the Gaussian process.




Fig. 2 Comparison of the posterior for (a) a Gaussian process and (b) a Student's-T process. The two processes start with the same prior mean and covariance and are updated with the same sample data. The posterior uncertainty of the Student's-T distribution is larger than the uncertainty in a GP, which compounded with STPs having larger outliers in general.

B. Illustrative Comparison

We illustrate the difference between the GP and STP processes. First, we compare draws from the prior distribution (before the function has been evaluated). We set both processes to have a zero mean function ($m(x) = 0$), and to use the isotropic squared exponential kernel function, (4). The STP is set with $\nu = 5.0$. Fig. 1 compares realizations of these two processes. The dark green line is the mean of the process at each (1-D) input, and the lighter green lines show the mean plus or minus two standard deviations. The lighter gray lines depict 300 realizations of the particular process, with Fig. 1a representing draws from the GP, and Fig. 1b showing realizations from the STP. Despite the fact that the two processes have the same mean and covariance, outliers are significantly more likely under the STP prior. The draws from the Gaussian process are mostly contained within the 2-standard-deviation error lines, and the most extreme deviations are not much outside it. The STP, on the other hand, has many more outliers, and there are several draws that far exceed the bounds, which would be very unlikely under a Gaussian prior.

We next compare the processes after observing data. Three evaluations from a (hypothetical) function are drawn, and the GP and STP priors are updated with these samples using the formulae in (5) and (14) respectively. The results of this update are shown in Fig. 2, with Fig. 2a showing the posterior of the GP, and Fig. 2b showing the posterior of the STP. We first observe that the STP is more likely to generate large outliers, just as under the prior. The GP only has significant outliers near the edge of the domain, while the STP has significant outliers not only on the edge of the domain, but also within the interior. Second, we see that the posterior uncertainty is generally larger for the STP than under the GP for this set of samples. These effects combine to make it much more likely to see a large outlier close to existing samples under an STP than a GP. As a result, if expected improvement is used to choose the next input to evaluate, an STP is more likely to choose an input close to an existing evaluation, while a GP is more likely to choose a location far from existing samples.

C. Marginal likelihood

It is often the case in Bayesian optimization that the kernel function contains of *hyperparameters* to be set, for example the bandwidth in (4). One common way to set the hyperparameters is to maximize the logarithm of the *marginal likelihood* of the data. The marginal likelihood of the STP can be found analytically from the data

$$\log(p(y|x)) = \gamma((\nu + d)/2) - \gamma(\nu/2) - \frac{d}{2}\log(\nu\pi) - \frac{1}{2}\log|\Sigma| - \frac{\nu + d}{2}\log\left(1 + \frac{y^T \Sigma^{-1} y}{\nu}\right) \quad (28)$$

where γ is the log of the Γ -function, and again, $\Sigma = \frac{\nu-2}{\nu}K$. The value of Σ , and thus the likelihood, depends on the hyperparameters of the kernel function. The hyperparameters in an STP can be chosen to maximize (28).

V. Numerical Experiments

We compare Student's-T processes with Gaussian processes on problems in Bayesian optimization. We first test the performance on synthetic benchmark functions commonly used in optimization, and then we compare the performance on an aerostructural optimization testbed problem.

A. Synthetic functions

The first synthetic function is the Rosenbrock function [18], given by

$$f(x_1, x_2) = (1 - x_1)^2 + 100(x_2 - x_1^2)^2 \\ x_1, x_2 \in [-3, 3] \quad (29)$$

The Rosenbrock function has a single local minimum of $f(x) = 0$ at $x = 0$. The second test function is the six-hump camel function [18], given by

$$f(x_1, x_2) = (4 - 2.1x_1^2 + \frac{x_1^4}{3})x_1^2 + x_1x_2 + (-4 + 4x_2^2)x_2^2 \\ x_1 \in [-3, 3], x_2 \in [-2, 2] \quad (30)$$

The six-hump camel function has six local minima. Two of these minima are global minima, which have a value of -1.0316 at $x = (0.0898, -0.7162)$ and $x = (-0.0898, 0.7162)$.

We compare the performance of a Gaussian Process, a Student's-T process with $\nu = 5$, and a Student's-T process with $\nu = 11$. All processes were set to have a mean function of 0, and to use the squared isotropic kernel function. The Bayesian optimization procedure was carried out identically for all of the processes. First, an initial set of 20 input locations were generated using Latin hypercube, and the objective function was evaluated at each location. The input and output data are scaled to have a mean of 0 and a variance of 1 in each dimension. Using this data, the best value of the kernel bandwidth parameter is found using a two-step grid search. Specifically, the marginal likelihood of the data as a function of the log of the bandwidth parameter is computed for 11 evenly spaced locations in $[-3, 3]$. The maximum of this initial grid search is used to refine a second grid search, again using 11 evenly spaced values. Then, the optimization procedure is run for 100 steps. At each step, the next input evaluated is chosen as the one with the greatest expected improvement. This location is found by first searching a grid with 101 locations in each dimension, and by using the grid location as the initial location for a local optimization. At every 10 steps, the evaluated inputs and outputs are re-normalized, and a new optimal bandwidth is found using the previously described procedure. The optimization is run until either 100 function evaluations have occurred, or until the global optimum is found to within 10^{-4} . The entire optimization procedure was repeated 100 different times (with different initial seedings) to find the average performance of the processes on the particular optimization problem. Note that for each individual optimization run, all processes begin with the same initial set of function evaluations.

Fig. 3 compares the performance on the Rosenbrock test case for the three processes. The y-axis shows the log of the difference between the current best found value and the global optimum y^* , with difference capped at a minimum of 10^{-4} . The x axis shows how the average of this value changes as a function of the optimization step. The shaded region shows the inner quartiles of performance over the 100 runs, and the solid line depicts the median performance. Similarly, Fig. 4 shows the performance on the six-hump camel test case. It can be seen that while all processes find a significantly better value than the initial seeding, the both of the STPs significantly outperform the GP. The STPs both find a better

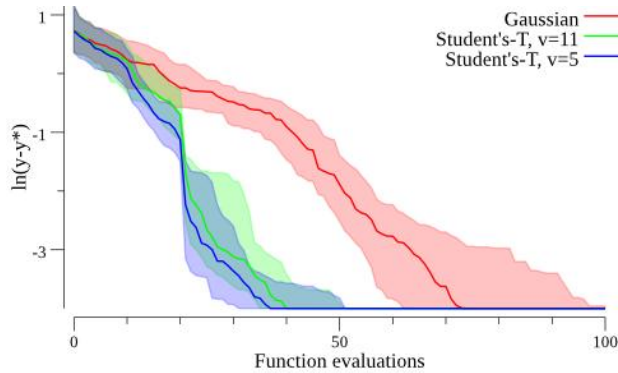


Fig. 3 A comparison of the GP against two STPs with different values of ν for the Rosenbrock test case. The y-axis shows $\log_{10}(\hat{y} - y^*)$ as a function of the optimization step. The solid line shows the median performance, and the shaded region covers the inner quartiles. Both STPs significantly outperform the GP.

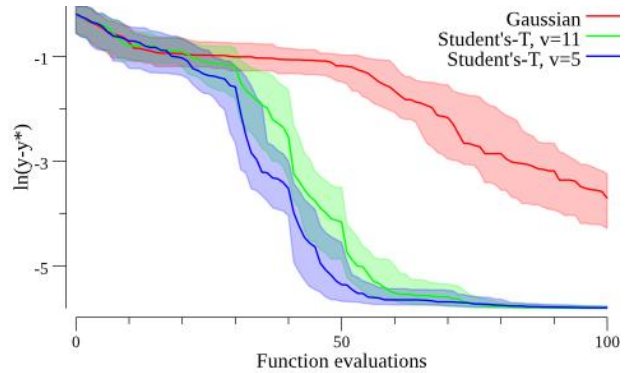


Fig. 4 A comparison of the GP against two STPs with different values of ν for the six-hump camel test case. The y-axis shows $\log_{10}(\hat{y} - y^*)$ as a function of the optimization step. The solid line shows the median performance, and the shaded region covers the inner quartiles. Both STPs significantly outperform the GP.

final value, and hone in on the optimum earlier in the optimization. Note also that the STPs are significantly more robust. Nearly all of the runs of the STP runs find the global minimum within the allotted budget, while in many cases the GP fails to find the best input. This is especially true for the six-hump camel function which contains many local minima, exactly the case where Bayesian optimization is most appropriate. The STP with $\nu = 5$ seems to outperform the STP with $\nu = 11$, though the differences are more minor compared with the differences with the GP.

B. Aerostructural optimization

The last test case is to find the optimal wing design for a coupled aerostructural problem, shown in Fig. 5. The coupled solver code is the OpenAeroStruct package [19], which is implemented on top of the OpenMDAO framework [20]. This test case has a 7-dimensional input, a univariate output objective, and two non-linear constraints. The input parameters are a) the angle of attack of the wing, constrained to be in $[-10, 10]$ degrees b) three wing thickness parameters constrained to be in $[0.01, 0.5]$ c) three wing twist parameters, constrained to be in $[-15, 15]$ degrees. The objective is to minimize the fuel burn of a simulated mission, and the wing twist and thickness are constrained so that a) the design is physically realizable (the structure does not intersect with itself) and b) the aerodynamic forces do not cause the wing to break. We transform this constrained optimization problem into the following unconstrained problem for demonstration purposes. If the constraints are not violated, the objective is simply taken as the fuel burn. If the constraints are violated, then the objective is taken to be $180000(1 + \xi)$, where ξ is the constraint violation (the value of 180000 was chosen to be larger than the fuel burn of a valid design)*. The optimization is performed almost identically as described above for the synthetic functions, except in a 7 dimensional input space a full grid search for the best EI is too computationally intensive. Instead, 10,000 test candidates are generated from Latin-hypercube sampling, and the best of those locations is used as the input to the local optimization step.

Fig. 6 compares the performance of the GP against the two STPs for the aerostructural design problem. Here, the plot shows the log of the current best value (rather than the normalized best value) since the global optimum is unknown. While the difference isn't as stark as the analytic cases, it is still clear that the STPs outperform the GP on this test case. The median performance is consistently better as the optimization progresses, and as the optimization continues, the 75th percentile of the STPs is close to outperforming the 25th percentile of the GP. The difference between the performance between the two STPs is insignificant for this problem.

*For some input conditions, the design constraints are satisfied, but a negative fuel burn is returned. This is clearly non-physical, and in this case the objective value is taken to be 300000. There are a small number of inputs for which the simulation crashes, and in this case the objective is taken to be 400000.




Fig. 5 A depiction of the OpenAeroStruct problem containing wing planform and internal structure




Fig. 6 A comparison of the GP against two STPs with different values of ν for the OpenAeroStruct test case. The y-axis shows $\log_{10}(\hat{y})$ as a function of the optimization step. The solid line shows the median performance, and the shaded region covers the inner quartiles. The STPs outperform the GP, though the two STPs behave nearly identically.

VI. Conclusion

This work has presented the Student's-T process, a prior over functions based on the multivariate Student's-T distribution. The STP has similar desirable properties to a Gaussian process, in that it has a simple expression for marginal distributions, and it has an analytic Bayesian update rule when new samples are observed. The STP also has some significant advantages over a Gaussian process. First, outliers are likely to occur under an STP with a small number of degrees of freedom. Second, the posterior covariance adjusts depending on the actual function values observed (and not just their locations), increasing if the samples vary by more than expected, or decreasing if they vary by less than expected. This work also presented an analytic expression for the expected improvement under a Student's-T distribution, and showed how to set kernel hyperparameters using marginal likelihood. Finally, we presented numerical simulation results that show STPs outperform GPs on several synthetic benchmarks as well as an aerosturctural design optimization problem.

It is known that STPs cannot be better than GPs for *every* possible design problem. However, it seems that the STP prior may be more naturally suited to problems in aerospace optimization. Aerospace design problems often feature smooth regions punctuated by (near) discontinuities, for example the transition to stall, or the failure to meet constraints. These deviations would be considered extremely unlikely according to a Gaussian prior, but not so under a Student's-T prior. Given their advantages, and no obvious disadvantages, it seems natural to “upgrade” from Gaussian processes to Student's-T processes in many aerospace design applications.

VII. Acknowledgements

This work was made possible through the support of AFOSR MURI on multi-information sources of multi-physics systems under Award Number FA9550-15-1-0038. We would also like to thank the Santa Fe Institute for support of this research.

References

- [1] Slotnick, J., Khodadoust, A., Alonso, J., Darmofal, D., Gropp, W., Lurie, E., and Mavriplis, D., “CFD vision 2030 study: a path to revolutionary computational aerosciences,” 2014.
- [2] Nelder, J. A., and Mead, R., “A simplex method for function minimization,” *The computer journal*, Vol. 7, No. 4, 1965, pp. 308–313.
- [3] Nocedal, J., and Wright, S., *Numerical optimization*, Springer Science & Business Media, 2006.
- [4] Deb, K., Agrawal, S., Pratap, A., and Meyarivan, T., “A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: NSGA-II,” *International Conference on Parallel Problem Solving From Nature*, Springer, 2000, pp. 849–858.

- [5] Van Laarhoven, P. J., and Aarts, E. H., “Simulated annealing,” *Simulated Annealing: Theory and Applications*, Springer, 1987, pp. 7–15.
- [6] Jones, D. R., Schonlau, M., and Welch, W. J., “Efficient global optimization of expensive black-box functions,” *Journal of Global optimization*, Vol. 13, No. 4, 1998, pp. 455–492.
- [7] Lukaczyk, T., Taylor, T., Palacios, F., and Alonso, J., “Managing gradient inaccuracies while enhancing optimal shape design methods,” *51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, 2013, p. 1042.
- [8] Lam, R., Allaire, D. L., and Willcox, K. E., “Multifidelity optimization using statistical surrogate modeling for non-hierarchical information sources,” *56th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, 2015, p. 0143.
- [9] Frazier, P., Powell, W., and Dayanik, S., “The knowledge-gradient policy for correlated normal beliefs,” *INFORMS journal on Computing*, Vol. 21, No. 4, 2009, pp. 599–613.
- [10] Rasmussen, C. E., “Gaussian processes for machine learning,” 2006.
- [11] Genz, A., and Bretz, F., *Computation of multivariate normal and t probabilities*, Vol. 195, Springer Science & Business Media, 2009.
- [12] Yu, S., Tresp, V., and Yu, K., “Robust multi-task learning with t-processes,” *Proceedings of the 24th international conference on Machine learning*, ACM, 2007, pp. 1103–1110.
- [13] Archambeau, C., and Bach, F., “Multiple Gaussian process models,” *arXiv preprint arXiv:1110.5238*, 2011.
- [14] Shah, A., Wilson, A. G., and Ghahramani, Z., “Student-t Processes as Alternatives to Gaussian Processes.” *AISTATS*, 2014, pp. 877–885.
- [15] Shah, A., Wilson, A. G., and Ghahramani, Z., “Bayesian Optimization using Student-t Processes,” ????
- [16] Roth, M., *On the multivariate t distribution*, Linköping University Electronic Press, 2012.
- [17] Slotani, M., “Tolerance regions for a multivariate normal population,” *Annals of the Institute of Statistical Mathematics*, Vol. 16, No. 1, 1964, pp. 135–153.
- [18] Molga, M., and Smutnicki, C., “Test functions for optimization needs,” *Test functions for optimization needs*, 2005, p. 101.
- [19] Jasa, J. P., Hwang, J. T., and Martins, J. R. R. A., “Open-source coupled aerostructural optimization using Python,” *Structural and Multidisciplinary Optimization*, 2018. (Submitted). Code retrieved from <http://github.com/mdolab/OpenAeroStruct>, last commit Nov 8, 2017.
- [20] Gray, J., Moore, K. T., and Naylor, B. A., “OpenMDAO: An open source framework for multidisciplinary analysis and optimization,” *AIAA/ISSMO Multidisciplinary Analysis Optimization Conference Proceedings*, Vol. 5, 2010.